FortiSwitchOS Administration Guide—Standalone Mode

VERSION 3.6.4

FORTINET DOCUMENT LIBRARY

http://docs.fortinet.com

FORTINET VIDEO GUIDE

http://video.fortinet.com

FORTINET BLOG

https://blog.fortinet.com

CUSTOMER SERVICE & SUPPORT

https://support.fortinet.com

http://cookbook.fortinet.com/how-to-work-with-fortinet-support/

FORTIGATE COOKBOOK

http://cookbook.fortinet.com

FORTINET TRAINING SERVICES

http://www.fortinet.com/training

FORTIGUARD CENTER

http://www.fortiguard.com

FORTICAST

http://forticast.fortinet.com

END USER LICENSE AGREEMENT

http://www.fortinet.com/doc/legal/EULA.pdf

FORTINET PRIVACY POLICY

https://www.fortinet.com/corporate/about-us/privacy.html

FEEDBACK

Email: techdocs@fortinet.com

Tuesday, January 16, 2018

FortiSwitchOS 3.6.4 Administration Guide Standalone Mode

Change log

Date	Change Description
January 12, 2018	Initial release.
January 16, 2018	Updated “Marking.”

TABLE OF CONTENTS

Introduction

This guide provides information about configuring a FortiSwitch unit in standalone mode. In standalone mode, you manage the FortiSwitch by connecting directly to the unit, either using the Web-based manager (also known as the GUI) or the CLI.

If you will be managing your FortiSwitch using a FortiGate unit, please see the following guide: Managing a FortiSwitch with a FortiGate.

This chapter covers the following topics:

Supported models

This guide is for all FortiSwitch models that are supported by FortiSwitchOS, which includes all of the D-series models.

What's new in FortiSwitchOS 3.6.4

Release 3.6.4 provides the following new features:

- Unicast hashing using the source port
- STP supported in MCLAGs
- QoS marking
- MAB reauthentication disabled by default
- Cut-through switching mode for low latency
- Control of how often the temperature and PoE alerts are generated
- Querier for IGMP snooping
- Logging of MAC address learning limit violations
- Persistent (sticky) MAC addresses and static MAC addresses saved to the same table
- Control of forwarding reserved multicast packets and forwarding IPv6 neighbor-discovery packets to the CPU for 124D, 124D-POE, 200 Series, and 400 Series
- New REST API endpoints

Refer to Feature matrix: FortiSwitchOS 3.6 on page 12 for details about the features supported on each FortiSwitch model.

Feature matrix: FortiSwitchOS 3.6

The following table lists the switch features in Release 3.6 that are supported on each series of switch models. All features are available in Release 3.6.0, unless otherwise stated.

Feature	GUI supported	108D-POE 112D-POE	1xxE	124D 124D-POE 200 Series 400 Series	500 Series	1024D 1048D	3032D
Link aggregation group size (maximum number of ports) (See Note 2.)	✓	8	8	8	24/48	24/48	24 (3.5.0) 64 (3.5.1)
Auto module max speed detection and notification	✓	—	—	—	✓	✓	—
IP conflict detection and notification	—	✓	✓	✓	✓	✓	✓
MAC-IP binding	✓	—	—	—	✓	✓	✓
Static BFD	—	—	—	—	—	✓	✓
Hardware-based ECMP	—	—	—	—	✓	✓	✓
Private VLANs	✓	—	—	✓	✓	✓	✓
Loop guard	✓	✓	✓	✓	✓	✓	✓
LAG min-max-bundle	—	✓	✓	✓	✓	✓	✓
sFlow	✓	✓	—	✓	✓	✓	✓
Storm control	✓	✓	✓	✓	✓	✓	✓
ACL	—	—	—	✓	✓	✓	✓
Static L3/hardware-based routing	✓	—	—	✓	✓	✓	✓
Software routing only	✓	✓	✓	—	—	—	—
CPLD software upgrade support for OS	—	—	—	—	—	✓	—

Feature	GUI supported	108D-POE 112D-POE	1xxE	124D 124D-POE 200 Series 400 Series	500 Series	1024D 1048D	3032D
PoE-pre-standard detection (See Note 1.)	✓	✓	FS-1xxE POE	✓	✓	—	—
VLAN tag by ACL	—	—	—	✓	✓	✓	✓
ACL redirect to mirror destination as trunk/LAG	—	—	—	✓	✓	✓	✓
MAC/IP/protocol-based VLAN assignment	✓	✓	—	✓	✓	✓	✓
802.1x port mode	✓	✓	✓	✓	✓	✓	✓
802.1x MAC-based security mode	✓	✓	✓	✓	✓	✓	✓
User-based (802.1x) VLAN assignment	✓	✓	—	✓	✓	✓	✓
Virtual wire	✓	—	—	✓	✓	✓	✓
HTTP REST APIs for configuration and monitoring	—	✓	✓	✓	✓	✓	✓
Split port	—	—	—	—	✓	—	✓
IGMP snooping	—	—	—	✓	✓	✓	✓
Per-port max for learned MACs	—	—	✓	✓	✓	—	—
802.1p support, including priority queuing trunk and WRED (release 3.5.1)	—	—	—	✓	✓	✓	✓
DHCP snooping	—	—	—	✓	✓	✓	✓
LLDP-MED	—	✓	✓	✓	✓	✓	✓
DHCP relay feature	—	—	✓	✓	✓	✓	✓
Support for switch SNMP OID	—	✓	✓	✓	✓	✓	✓

Feature	GUI supported	108D-POE 112D-POE	1xxE	124D 124D-POE 200 Series 400 Series	500 Series	1024D 1048D	3032D
Access VLANs (See Note 5.)	—	—	—	✓	✓	✓	✓
802.1x enhancements, including MAB (release 3.5.1)	✓	✓	✓	✓	✓	✓	✓
Multi-stage load balancing (release 3.5.1)	—	—	—	—	—	✓	✓
MCLAG (multichassis link aggregation)(release 3.6.0)	—	—	—	✓ (not on 124D/124D-POE)	✓	✓	✓
Dynamic layer-3 protocols (OSPF, RIP, and VRRP) (release 3.6.0) (See Note 3.)	✓	—	—	✓ (not on 124D/124D-POE)	✓	✓	✓
Dynamic ARP inspection (release 3.6.0)	—	—	—	✓	✓	✓	✓
Firmware image rotation (dual-firmware image support) (release 3.6.0)	—	✓ (not on 108D-POE)	—	✓	✓	✓	✓
TDR (time-domain reflectometer)/cable diagnostics support (release 3.6.0)	✓	—	—	✓	✓	✓	✓
MAC learning limit (release 3.6.0) (See Note 4.)	—	—	✓	✓	✓	—	—
Sticky MAC on switch interfaces (release 3.6.0)	—	—	—	✓	✓	✓	✓
PoE modes support: first come, first served or priority based (PoE models) (release 3.6.0)	—	✓	FS-1xxE POE	✓	✓	—	✓

Feature	GUI supported	108D-POE 112D-POE	1xxE	124D 124D-POE 200 Series 400 Series	500 Series	1024D 1048D	3032D
ACL: egress mask action support (release 3.6.0)	—	—	—	✓	✓	✓	✓
Monitor system temperature (threshold configuration and SNMP trap support) (release 3.6.0)	—	✓	—	✓	✓	✓	✓
'forced-untagged' or 'force-tagged' setting on switch interfaces (release 3.6.0)	—	✓	—	✓	✓	✓	✓
Selective packet sampling to CPU (useful diagnostic tool) (release 3.6.0)	—	—	—	✓	✓	✓	3.6.1
Add CLI to show the details of port statistics (release 3.6.0)	—	✓	✓	✓	✓	✓	✓
Display progress (%) during firmware upgrade (release 3.6.0)	✓	✓	✓	✓	✓	✓	✓
STP root guard (release 3.6.2)	—	✓	✓	✓	✓	✓	✓
STP BPDU guard (release 3.6.2)	—	✓	✓	✓	✓	✓	✓
IGMP snooping: static multicast groups (release 3.6.2)	—	—	—	✓	✓	✓	✓
DHCP snooping: entry limit per port (release 3.6.2)	—	—	—	✓	✓	✓	✓
Network device detection (release 3.6.2)	—	—	—	✓	✓	✓	✓

Feature	GUI supported	108D-POE 112D-POE	1xxE	124D 124D-POE 200 Series 400 Series	500 Series	1024D 1048D	3032D
QoS queue counters (releases 3.6.2 and 3.6.3)	—	—	—	✓	✓	✓	✓
Support of the RADIUS accounting server (release 3.6.3)	—	✓	—	✓	✓	✓	✓
Support of RADIUS CoA and disconnect messages (release 3.6.3)	—	✓	—	✓	✓	✓	✓
802.1x authentication: EAP-TLS support (release 3.6.3)	—	✓	—	✓	✓	✓	✓
DHCP snooping: CLI for DHCP-snooping server database (release 3.6.3)	—	—	—	✓	✓	✓	✓
Unicast hashing (release 3.6.4)	—	—	—	✓	✓	✓	✓
STP supported in MCLAGs (release 3.6.4)	—	—	—	✓ (not on 124D/124D-POE)	✓	✓	✓
QoS marking (release 3.6.4)	—	—	—	✓	✓	✓	✓
MAB reauthentication disabled (release 3.6.4)	—	✓	—	✓	✓	✓	✓
Cut-through switching (release 3.6.4)	—	—	—	—	—	✓	✓
Control of temperature and PoE alerts (release 3.6.4)	—	✓	—	✓	✓	✓	✓
IGMP querier (release 3.6.4)	—	—	—	✓	✓	✓	✓

Feature	GUI supported	108D-POE 112D-POE	1xxE	124D 124D-POE 200 Series 400 Series	500 Series	1024D 1048D	3032D
Configuration of the QSFP low-power mode (release 3.6.4)	—	—	—	—	✓	1048D	✓
Learning limit violation log (release 3.6.4) (See Note 4.)	—	—	—	✓	✓	—	—
Sticky MAC addresses saved to static MAC table (release 3.6.4)	—	—	—	✓	✓	✓	✓
Enabling packet forwarding to CPU (release 3.6.4)	—	—	—	✓	—	—	—

Notes

1. PoE features are applicable only to the model numbers with a POE or FPOE suffix.
2. 24-port LAG is applicable to 524D, 524_FPOE, 1024D, and 3032D models. 48-port LAG is applicable to 548D, 548_FPOE, and 1048D models.
3. To use the dynamic layer-3 protocols, you must have an advanced features license.
4. The per-VLAN learning limit and per-trunk learning limit are not supported on dual-chip platforms (248 and 448 series).
5. Applicable to the FortiLink managed mode.

Before you begin

Before you start administrating your FortiSwitch unit, it is assumed that you have completed the initial configuration of the FortiSwitch unit, as outlined in the QuickStart Guide for your FortiSwitch model and have administrative access to the FortiSwitch unit's Web-based manager and CLI.

How this guide is organized

This guide is organized into the following chapters:

- Management ports - configuring the management ports.
- Configuring administrator tasks - configuring date and time, admin users, remote authentication servers.
- Configuring SNMP - allows you to monitor hardware on your network.
- Global system settings - the initial configuration of your FortiSwitch unit.
- Physical port settings - configuring the physical ports.

- Layer-2 interfaces - configuring layer-2 interfaces.
- VLANs and VLAN tagging - configuration and packet flow for VLAN-tagged and untagged packets.
- Spanning Tree Protocol - how to configure MSTP.
- Link aggregation groups - configuring link aggregation groups.
- MCLAG - configuring MCLAG.
- Multi-stage load balance - configuring multi-stage load balancing on a set of FortiGate units.
- LLDP-MED - how to configure LLDP-MED settings.
- MAC/IP/protocol-based VLANs - configuring MAC/IP/protocol-based VLANs.
- Mirroring - configuring port mirroring.
- Access control lists - configuring ACLs.
- Storm control - configuring storm control.
- DHCP snooping - configuring DHCP snooping.
- Dynamic ARP inspection - configuring dynamic ARP inspection.
- IGMP snooping - configuring IGMP snooping.
- Private VLANs - creation and management of private virtual local area networks (VLANs).
- QoS settings - how to configure QoS.
- sFlow - configuring sFlow.
- Feature licensing - about feature licenses.
- Layer-3 interfaces - configuring routed ports, routed VLAN interfaces, switch virtual interfaces, and features related to these interfaces.
- DHCP relay - configuring DHCP relay.
- OSPF routing - configuring OSPF routing.
- RIP routing - configuring RIP routing.
- VRRP - configuring VRRP.
- Users and user groups - configuring users and user groups.
- 802.1x authentication - configuring 802.1x authentication (to RADIUS servers).
- TACACS - configuring TACACS authentication.
- Troubleshooting and support
- Deployment scenario

Management ports

This chapter describes how to configure management ports on the FortiSwitch.

The following topics are covered:

Models without a dedicated management port

For FortiSwitch models without a dedicated management port, configure the internal interface as the management port.

Note: For FortiSwitch models without a dedicated management port, the internal interface has a default VLAN ID of 1.

Using the Web-based manager:

First start by editing the default **internal** interface's configuration.

1. Go to **System > Network > Interface**, click **internal** to select it, and then click **Edit**.

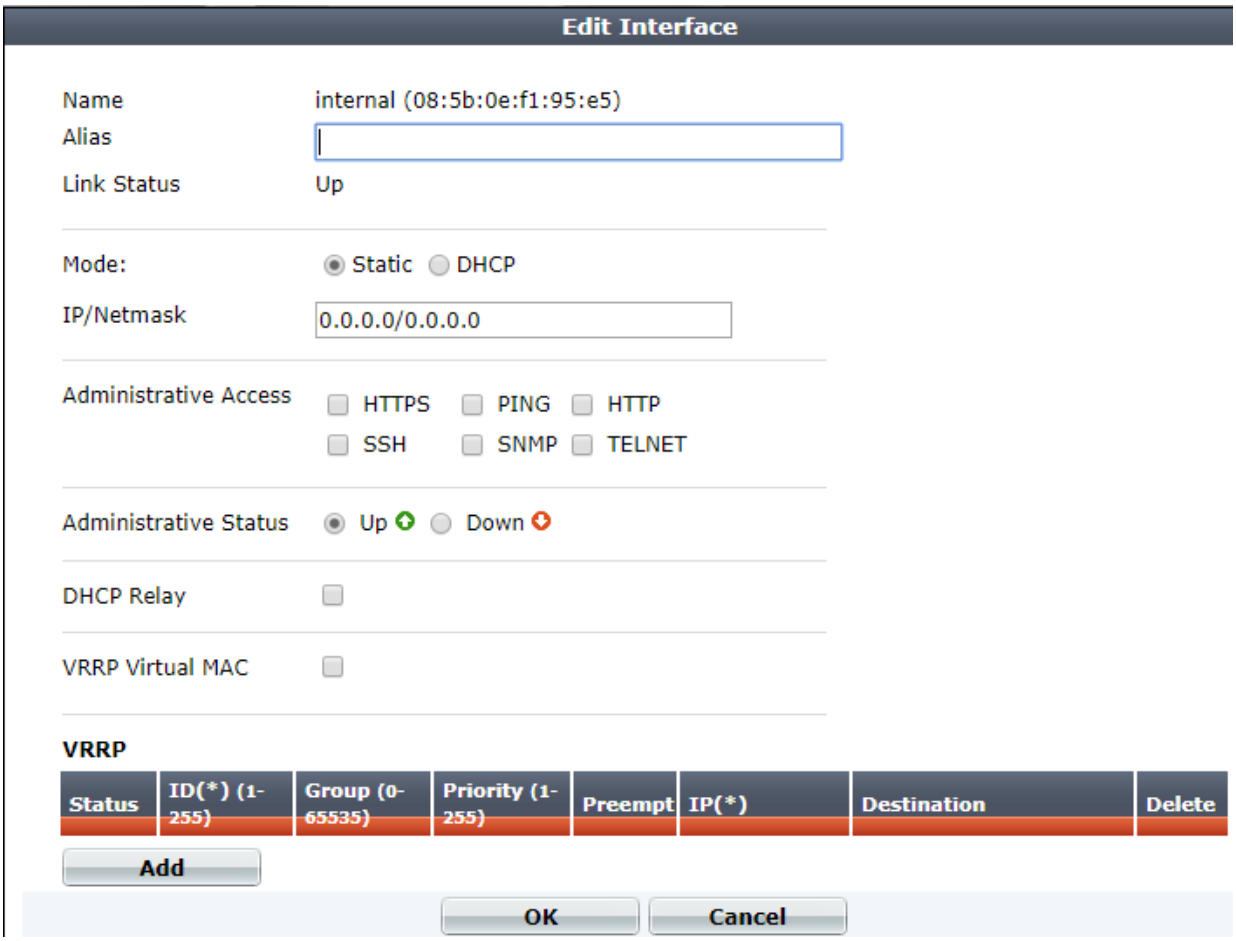

2. In the IP/Netmask field, enter the IP address and netmask.
3. Select the appropriate protocols to connect to the interface for administrative access.
4. Click **OK**.

Next, create a new interface to be used for management.

1. Go to **System > Network > Interface** and click **Create New** to create a management VLAN.

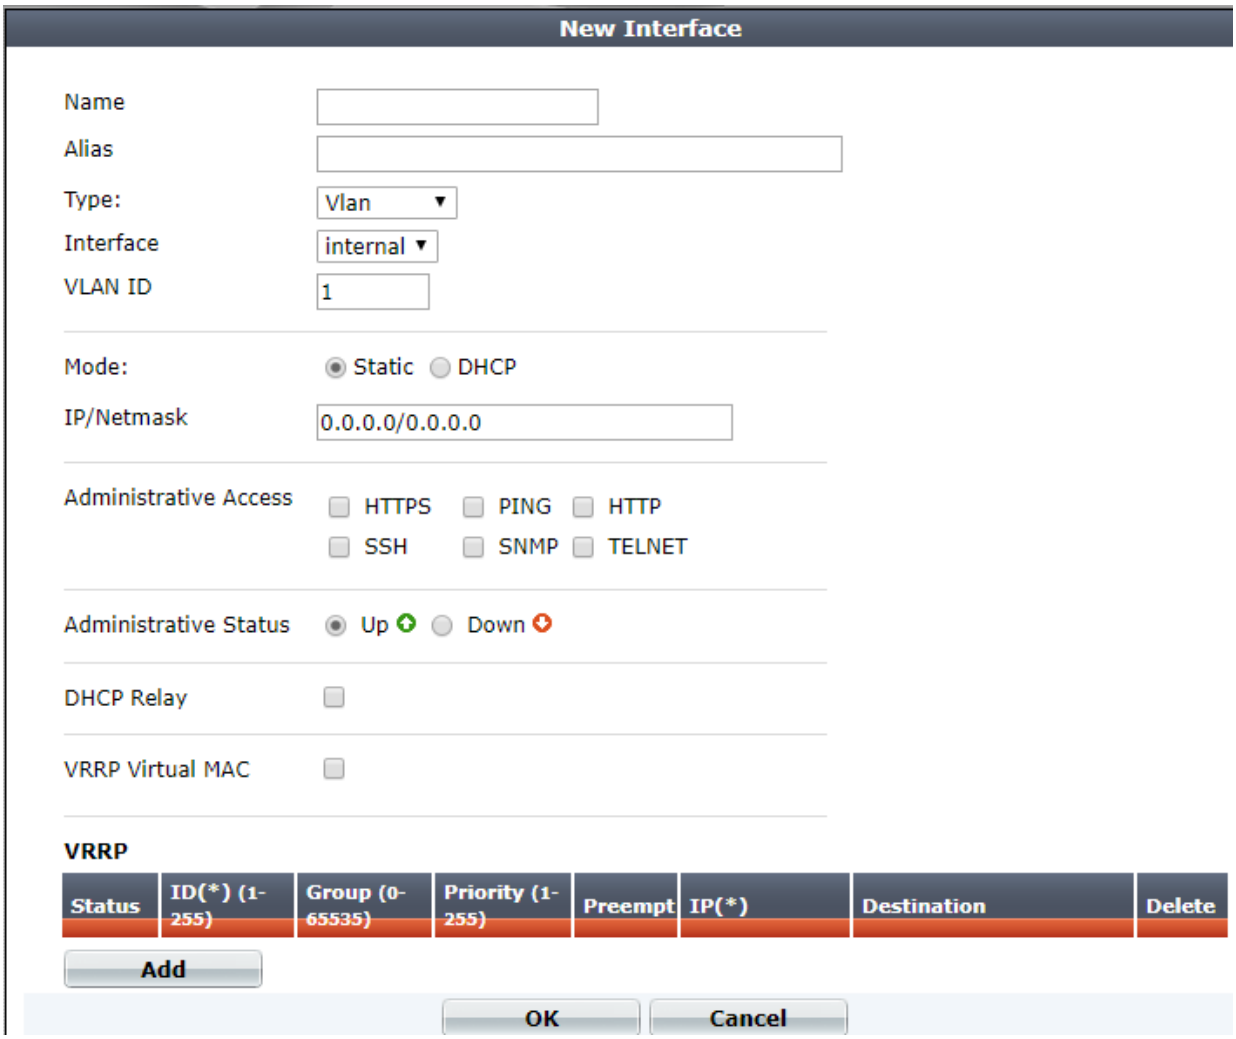

2. Give the interface an appropriate name.
3. Set **Interface** to **internal**.
4. Set a **VLAN ID**.
5. In the IP/Netmask field, enter the IP address and netmask.
6. Select the appropriate protocols to connect to the interface for administrative access.
7. Click **OK**.

Using the CLI:

```
config system interface
   edit internal
      set ip <address>
      set allowaccess <access_types>
      set type physical
   next
   edit <vlan name>
      set ip <address>
      set allowaccess <access_types>
      set interface internal
      set vlanid <VLAN id>
```

```
    end
end
```

Models with a dedicated management port

For FortiSwitch models with a dedicated management port, configure the IP address and allowed access types for the management port.

Note: For FortiSwitch models with a dedicated management port, the internal interface has a default VLAN id of 4094.

Using the Web-based manager:

1. Go to **System > Network > Interface** and click **mgmt** to select it, and then click **Edit**.

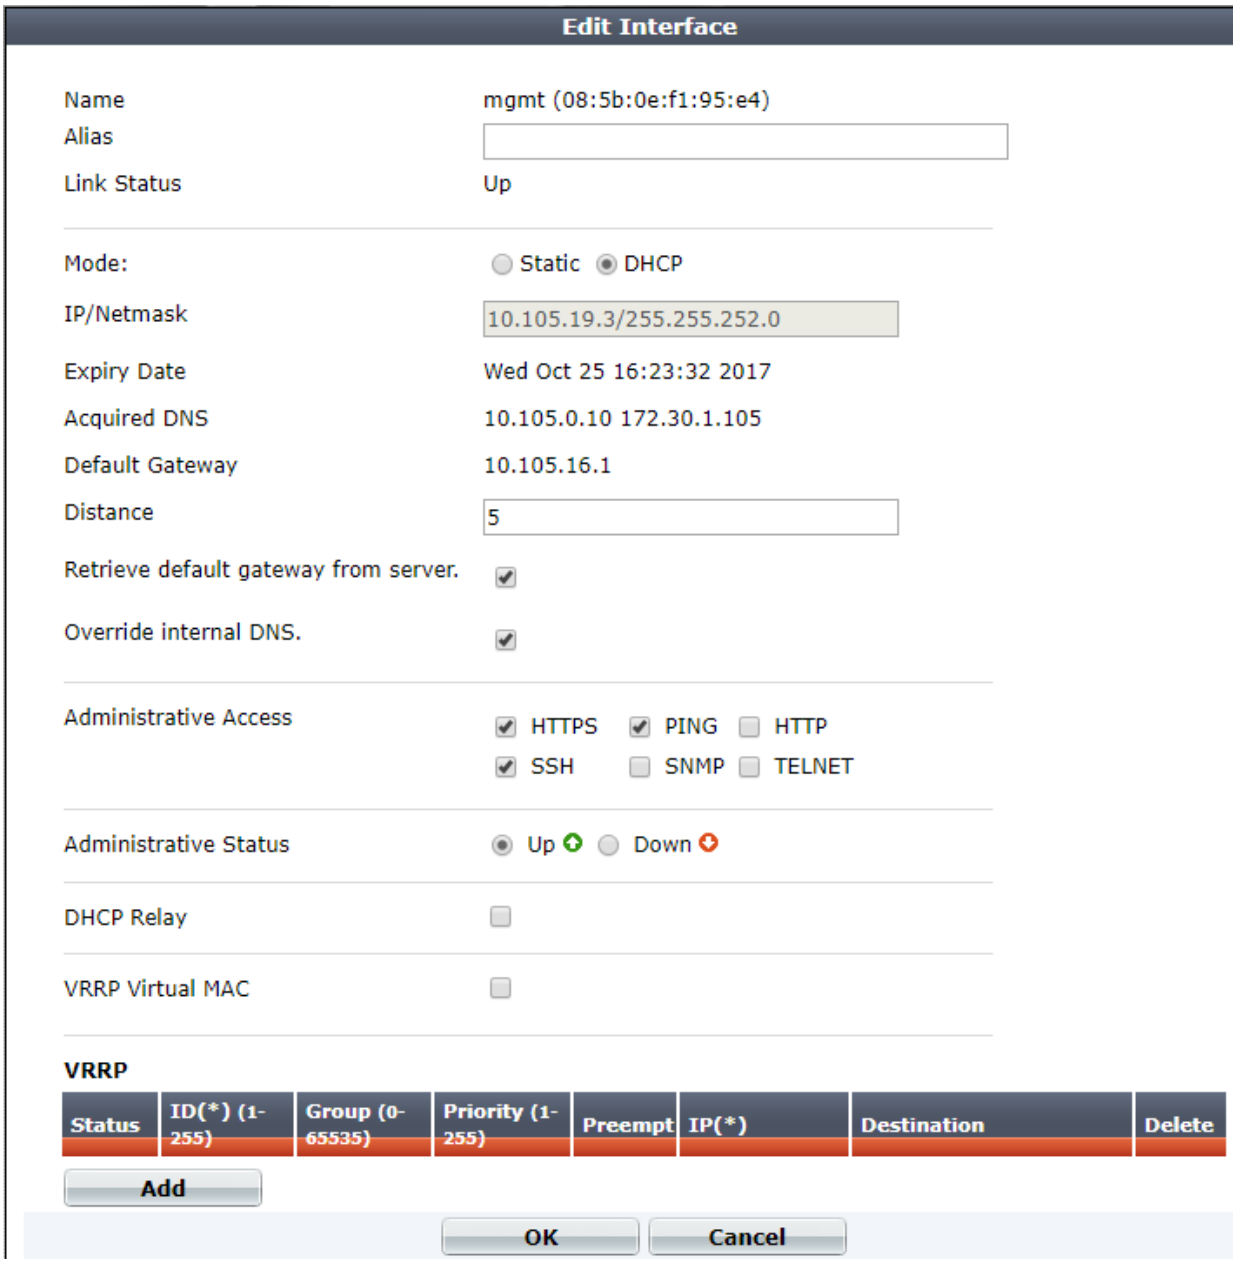

2. In the IP/Netmask field, enter the IP address and netmask.
3. Select the appropriate protocols to connect to the interface for administrative access.
4. Click **OK**.

Using the CLI:

```
config system interface
   edit mgmt
      set ip <address>
      set allowaccess <access_types>
      set type physical
   next
   edit internal
      set type physical
```

```
    end
end
```

Remote access to the management port

To provide remote access to the management port, configure a static route. Set the gateway address to the IP address of the router.

Using the Web-based manager:

1. Go to **Router > Router > Static Route** and click **Create New**.

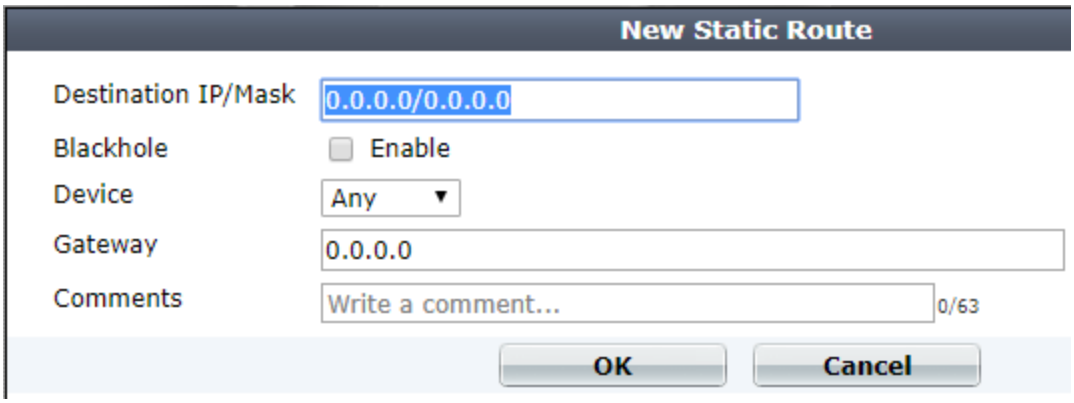

2. Set the device to **mgmt**.
3. Set the Gateway to the gateway router IP address.
4. Click **OK**.

Using the CLI:

```
config router static
   edit 1
      set device mgmt
      set gateway <router IP address>
   end
end
```

Example configurations

The following example configurations are for management ports, with the CLI syntax shown to create them.

In this example, the **internal** interface is used as an inbound management interface. Also, the FortiSwitch has a default VLAN across all physical ports and its internal port.

Using the internal interface of a FortiSwitch-108D-POE

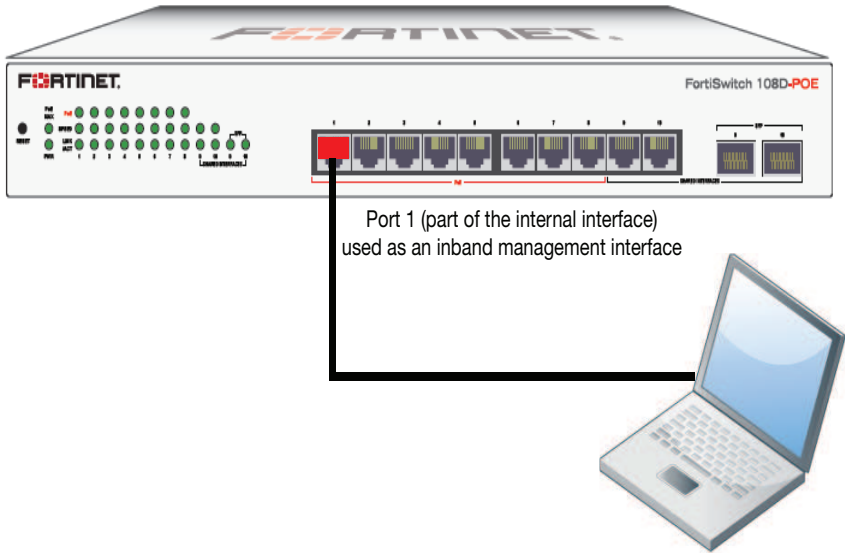

Syntax

```
config system interface
   edit internal
      set ip 192.168.1.99 255.255.255.0
      set allowaccess ping https http ssh
      set type physical
   end
end
```

In the example, an out-of-band management interface is used as the dedicated management port. You can configure the management port for local or remote access.

Out of band management on a FortiSwitch-1024D

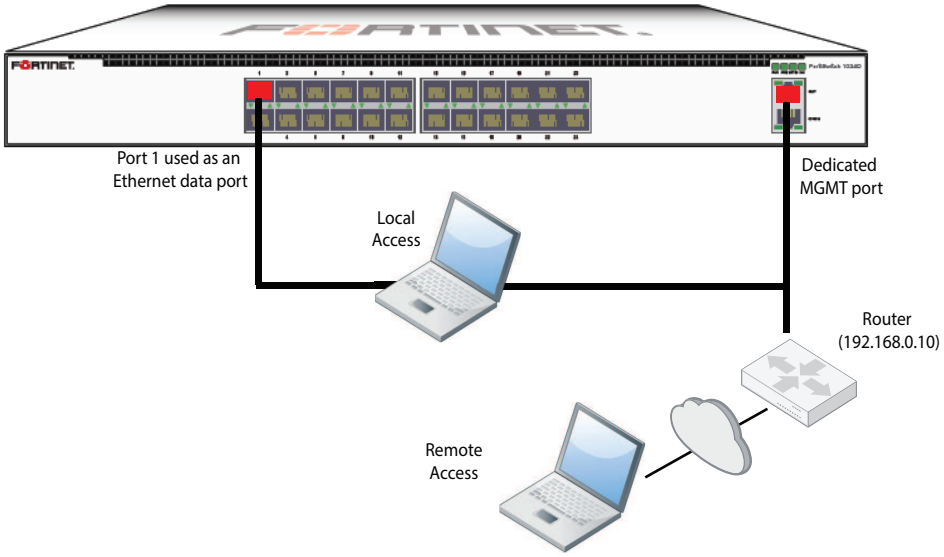

Option 1: management port with static IP

```
config system interface
   edit mgmt
      set ip 10.105.142.19 255.255.255.0
      set allowaccess ping https http ssh snmp telnet
      set type physical
   next
   edit internal
      set type physical
   end
end
// optional configuration to allow remote access to the management port

config router static
   edit 1
      set device mgmt
      set gateway 192.168.0.10
   end
end
```

Option 2: management port with IP assigned by DHCP

```
config system interface
   edit mgmt
      set mode dhcp
      set defaultgw enable // allows remote access
      set allowaccess ping https http ssh snmp telnet
      set type physical
```

```
    next
    edit internal
       set type physical
    end
end
```

Configuring administrator tasks

You can use the default “admin” account to configure administrator accounts, adjust system settings, upgrade firmware, create backup files, and configure security features.

This chapter covers the following topics:

Setting the time and date

For effective scheduling and logging, the system date and time must be accurate. You can either manually set the system date and time or configure the system to automatically keep its time correct by synchronizing with a Network Time Protocol (NTP) server.

The Network Time Protocol enables you to keep the system time synchronized with other network systems. This will also ensure that logs and other time-sensitive settings are correct.

To set the date and time

1. Go to **System** > **Dashboard** > **Status** and locate the **System Information** widget.

2. Beside **System Time,** select **Change**.

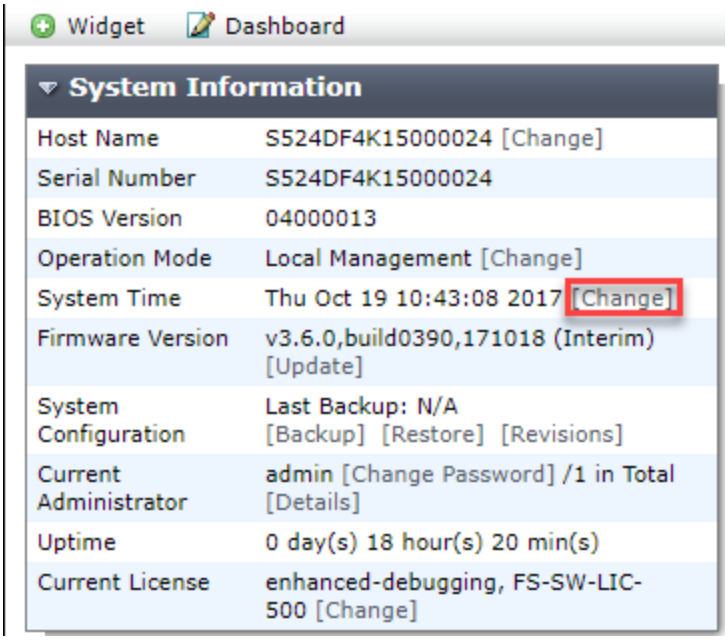

3. Select your **Time Zone**.
4. Either select **Set Time** and manually set the system date and time or select **Synchronize with NTP Server**. If you select synchronization, you can either use the default FortiGuard servers or specify a different server. You can also set the **Sync Interval**.
5. Click **OK**.

If you use an NTP server, you can identify a specific port/IP address for this self-originating traffic. The configuration is performed in the CLI with the command `set source-ip`. For example, you can set the source IP address of NTP to be on the DMZ1 port with an IP of 192.168.4.5:

```
config system ntp
   set ntpsyn enable
   set syncinterval 5
   set source-ip 192.168.4.5
end
```

Configuring the temperature sensor

If your FortiSwitch has a temperature sensor, you can set a warning and an alarm for when the system temperature reaches specified temperatures. When these thresholds are exceeded, a log message and SNMP trap are generated. The warning threshold must be lower than the alarm threshold.

Use the following commands to set warning and alarm thresholds:

```
config system snmp sysinfo
   set status enable
   set trap-temp-warning-threshold <temperature in degrees Celsius>
   set trap-temp-alarm-threshold <temperature in degrees Celsius>
end
```

By default, the FortiSwitch generates an alert (in the form of an SNMP trap and a SYSLOG entry) every 10 minutes when the temperature sensor exceeds its set threshold. You can change this interval with the following commands:

```
config system global
   set alert-interval <1-1440>
end
```

Configuring the PoE sensor

If your FortiSwitch has a PoE sensor, you can set an alarm for when the current power budget exceeds a specified percentage of the total power budget. When this threshold is exceeded, log messages and SNMP traps are generated. The default threshold is 80 percent.

Use the following commands to set the alarm threshold for the PoE sensor:

```
config switch global
   set poe-alarm-threshold <threshold (percent of total power budget) above which an alarm
       event is generated>
end
```

By default, the FortiSwitch generates an alert (in the form of an SNMP trap and a SYSLOG entry) every 10 minutes when the PoE sensor exceeds its set threshold. You can change this interval with the following commands:

```
config system global
   set alert-interval <1-1440>
end
```

Upgrading the firmware

Use these procedures to upgrade your FortiSwitch firmware.

Using the Web based manager:

You can upgrade the firmware from the dashboard or from the system configuration page.

To upgrade the firmware from the dashboard:

1. Go to **System** > **Dashboard** > **Status**.

2. Next to the **Firmware Version** field, click the **Update** link.

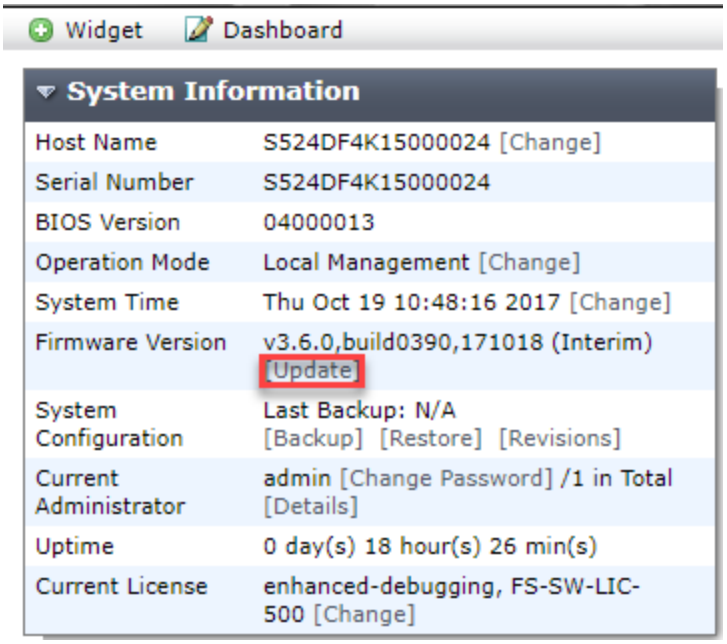

To upgrade the firmware from the system configuration page:

1. Go to **System** > **Config** > **Firmware**.
2. Click **Upgrade**.

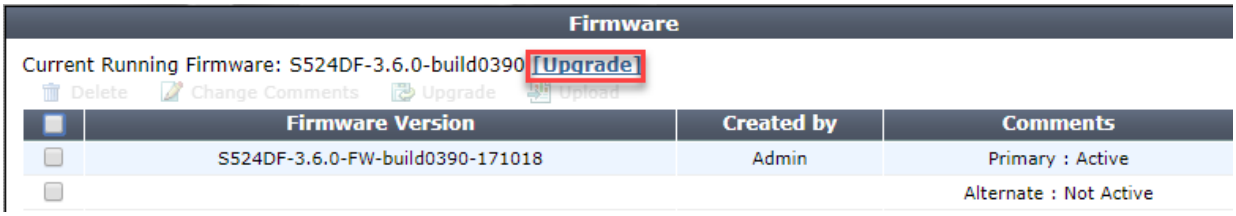

Using the CLI:

You can download a firmware image from an FTP server, from a FortiManager unit, or from a TFTP server. The FortiSwitch reboots and then loads the new firmware.

```
execute restore image ftp <filename_str> <server_ipv4[:port_int] | server_fqdn[:port_int]>
    [<username_str> <password_str>]
execute restore image management-station <version_int>
execute restore image tftp <filename_str> <server_ipv4>
```

The following example shows how to upload a configuration file from a TFTP server to the FortiSwitch and restart the FortiSwitch with this configuration. The name of the configuration file on the TFTP server is `backupconfig`. The IP address of the TFTP server is 192.168.1.23.

```
execute restore config tftp backupconfig 192.168.1.23
```

You can also load a firmware image from an FTP or TFTP server without restarting the FortiSwitch:

```
execute stage image ftp <string> <ftp server>[:ftp port]
execute stage image tftp <string> <ip>
```

Verifying image integrity

To verify the integrity of the images in the primary and secondary (if applicable) flash partitions, use the following commands:

```
execute verify image primary
execute verify image secondary
```

If the image is corrupted or missing, the command fails with a return code of -1.

For example:

```
execute verify image primary

Verifying the image in flash......100%
No issue found!

execute verify image secondary

Verifying the image in flash......100%
Bad/corrupted image found in flash!
Command fail. Return code -1
```

Restore or upgrade the BIOS

You can restore or upgrade the basic input/output system (BIOS) if needed.

CAUTION: Only restore or upgrade the BIOS if Customer Support recommends it.

To upgrade or restore the BIOS from the CLI:

```
execute restore bios tftp <filename_str> <server_ipv4[:port_int]>
```

For example:

```
execute restore bios tftp PPC/FS-3032D/04000009/FS3D323Z14000004.bin 10.105.2.201
```

The example downloads the BIOS file from the TFTP server at the specified IPv4 address.

NOTE: If the BIOS upgrade fails, do not restart the FortiSwitch. Instead, try the CLI command again. If repeating the CLI command does not work, the FortiSwitch might require a return merchandise authorization (RMA).

Setting the boot partition

You can specify the flash partition for the next reboot. The system can use the boot image from either the primary or the secondary flash partition:

```
execute set-next-reboot <primary|secondary>
```

If your FortiSwitch model has dual flash memory, you can use the primary and backup partitions for image rotation. By default, this feature is disabled.

```
config system global
   set image-rotation <enable | disable>
end
```

To list all of the flash partitions:

```
diagnose sys flash list
```

Backing up the system configuration

To back up the configuration from the dashboard:

1. Go to **System** > **Dashboard** > **Status**.
2. Next to the **System Configuration** field, click the **Backup** link.

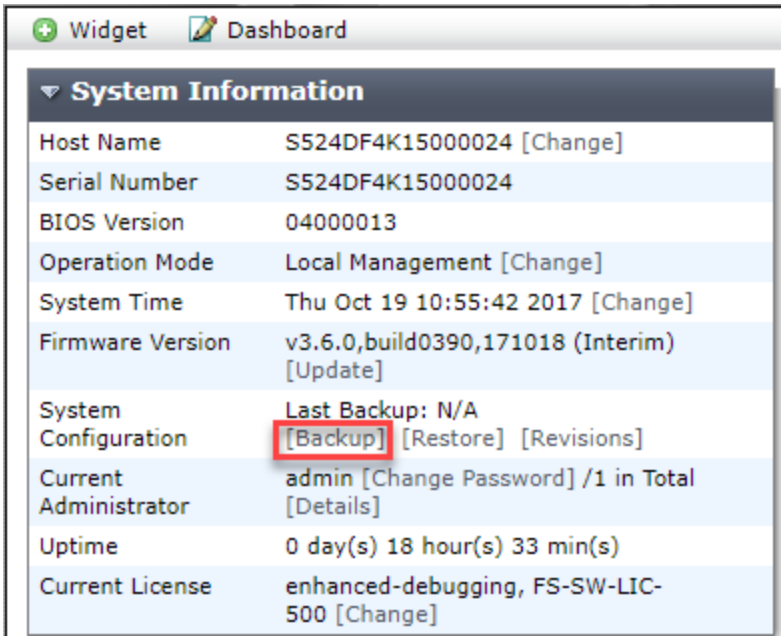

Remote authentication servers

If you are using remote authentication for administrators or users, you need to configure one of the following:

- RADIUS server
- TACACS+ server

RADIUS server

The information you need to configure the system to use a RADIUS server includes:

- the RADIUS server's domain name or IP address
- the RADIUS server's shared secret key

The default port for RADIUS traffic is 1812. Some RADIUS servers use port 1645. You can configure the FortiSwitch to use port 1645:

```
config system global
   set radius-port 1645
end
```

To configure RADIUS authentication with the Web-based manager:

1. Go to **System** > **Authentication** > **RADIUS Servers** and select **Create New**.

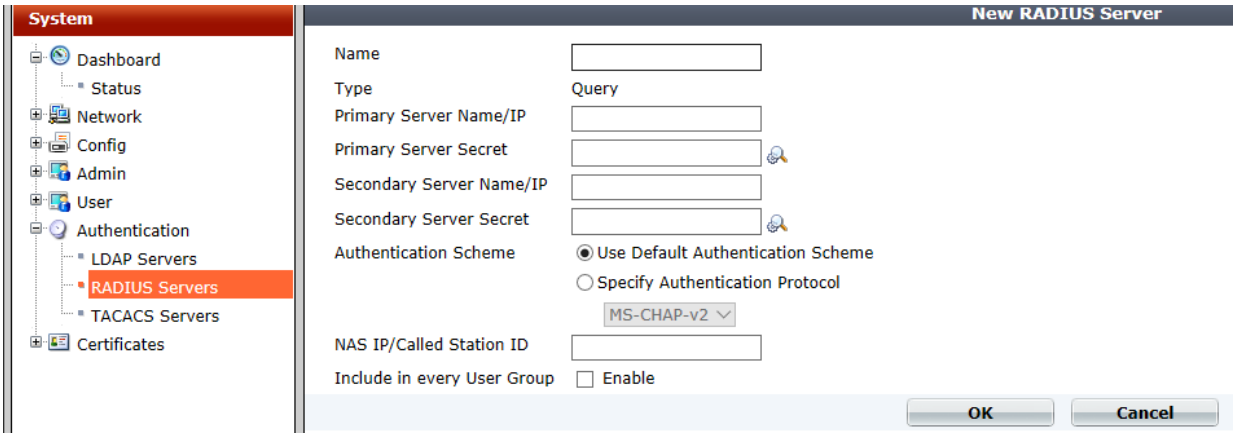

2. Enter the following information and click **OK**.

Field	Description
Name	Enter a name to identify the RADIUS server on the FortiSwitch.
Type	Select **Query** or **Dynamic Start**.
Primary Server Name/IP	Enter the domain name (such as fgt.exmaple.com) or the IP address of the RADIUS server.
Primary Server Secret	Enter the server secret key, such as radiusSecret. This key can be a maximum of 16 characters long. This value must match the secret on the RADIUS primary server.
Secondary Server Name/IP	Optionally enter the domain name (such as fgt.exmaple.com) or the IP address of the secondary RADIUS server.

Field	Description
Secondary Server Secret	Optionally, enter the secondary server secret key, such as radiusSecret2. This key can be a maximum of 16 characters long. This value must match the secret on the RADIUS secondary server.
Authentication Scheme	If you know the RADIUS server uses a specific authentication protocol, select it from the list. Otherwise select **Use Default Authentication Scheme**. The Default option will usually work.
NAS IP/ Called Station ID	Enter the IP address to be used as an attribute in RADIUS access requests. **NAS-IP-Address** is RADIUS setting or IP address of the FortiSwitch interface used to talk to the RADIUS server, if not configured. **Called Station ID** is the same value as NAS-IP Address but in text format.
Include in every User Group	When enabled, this RADIUS server will automatically be included in all user groups. This option is useful if all users will be authenticating with the remote RADIUS server.

To configure the FortiSwitch for RADIUS authentication, see 802.1x authentication on page 170.

TACACS+ server

TACACS+ is a remote authentication protocol that provides access control for routers, network access servers, and other networked computing devices using one or more centralized servers. TACACS+ allows a client to accept a user name and password and send a query to a TACACS+ authentication server. The server host determines whether to accept or deny the request and sends a response back that allows or denies the user access to the network.

TACACS+ offers fully encrypted packet bodies and supports both IP and AppleTalk protocols. TACACS+ uses TCP port 49, which is seen as more reliable than RADIUS's UDP protocol.

To configure TACACS+ authentication using the Web-based manager:

1. Go to **System** > **Authentication** > **TACACS Servers** and select **Create New**.

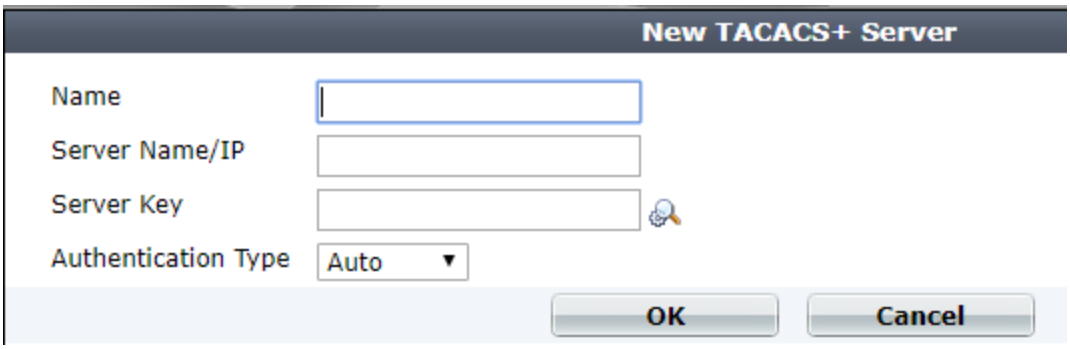

2. Enter the following information and click **OK**.

Field	Description
Name	Enter a name to identify the TACACS server on the FortiSwitch.
Server Name/IP	Enter the domain name (such as fgt.exmaple.com) or the IP address of the TACACS server.
Server Key	Enter the server key for the TACACS server.
Authentication Type	Select the authentication type to use for the TACACS+ server. **Auto** tries PAP, MSCHAP, and CHAP (in that order).

To configure the FortiSwitch for TACACS+ authentication, see TACACS on page 195.

Configuring system administrators

In addition to the default “admin” account, you might want to set up other administrators with different levels of system access.

This section covers the following topics:

Administrator profiles

Administer profiles define what the administrator user can do when logged into the FortiSwitch. When you set up an administrator user account, you also assign an administrator profile, which dictates what the administrator

user will see. Depending on the nature of the administrator's work, access level, or seniority, you can allow them to view and configure as much, or as little, as required.

The super_admin administrator is the administrative account that the primary administrator should have to log into the FortiSwitch. The profile cannot be deleted or modified to ensure there is always a method to administer the FortiSwitch. This user profile has access to all components of the system, including the ability to add and remove other system administrators. For some administrative functions, such as backing up and restoring the configuration using SCP, super_admin access is required.

Creating administrator profiles

To configure administrator profiles go to **System** > **Admin Profiles**. You can only assign one profile to each administrator user.

On the **New Admin Profile** page, you define the components of FortiSwitch that will be available to view and/or edit. For example, if you configure a profile so that the administrator can only access System Configuration, this admin will not be able to change Network settings.

Using the Web-based manager:

1. Go to **System > Admin > Admin Profile** and click **Create New**.

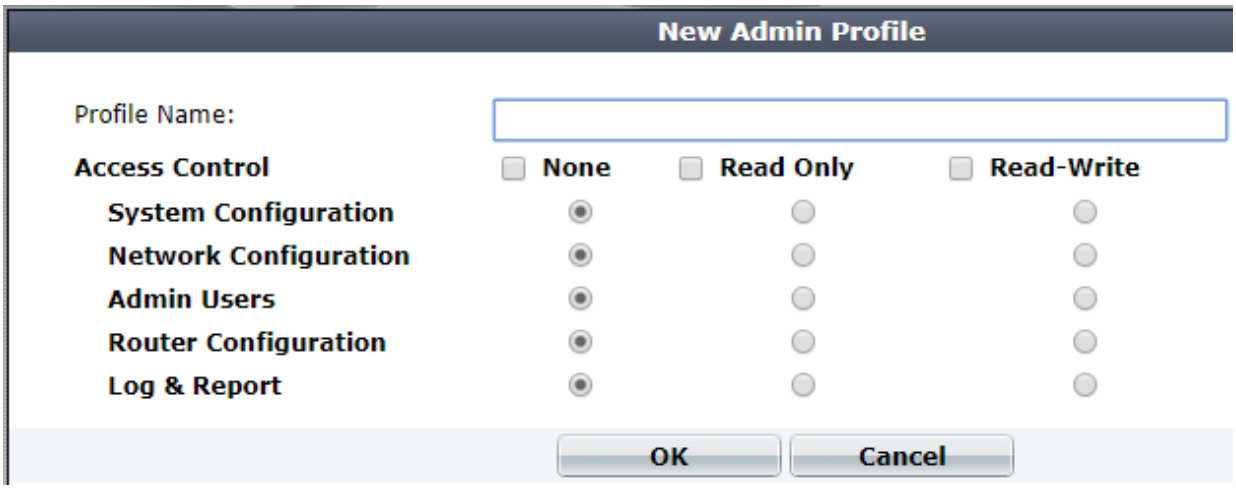

2. Give the profile an appropriate name.
3. Set **Access Control** as desired, choosing between **None**, **Read Only**, or **Read-Write.**
4. Click **OK**.

Using the CLI:

```
config system accprofile
   edit <name>
      set admingrp {none | read | read-write}
      set loggrp {none | read | read-write}
      set netgrp {none | read | read-write}
      set routegrp {none | read | read-write}
      set sysgrp {none | read | read-write}
   end
end
```

Adding administrators

Only the default “admin” account can create a new administrator account. If required, you can add an additional account with read-write access control to add new administrator accounts.

If you log in with an administrator account that does not have the super_admin admin profile, the administrators list will show only the administrators for the current virtual domain.

When adding administrators, you are setting up the administrator’s user account. An administrator account comprises an administrator’s basic settings as well as their access profile. The access profile is a definition of what the administrator is capable of viewing and editing.

Follow one of these procedures to add an administrator.

Using the Web-based manager:

1. Go to **System** > **Admin** > **Administrators**.
2. Click **Create New**.

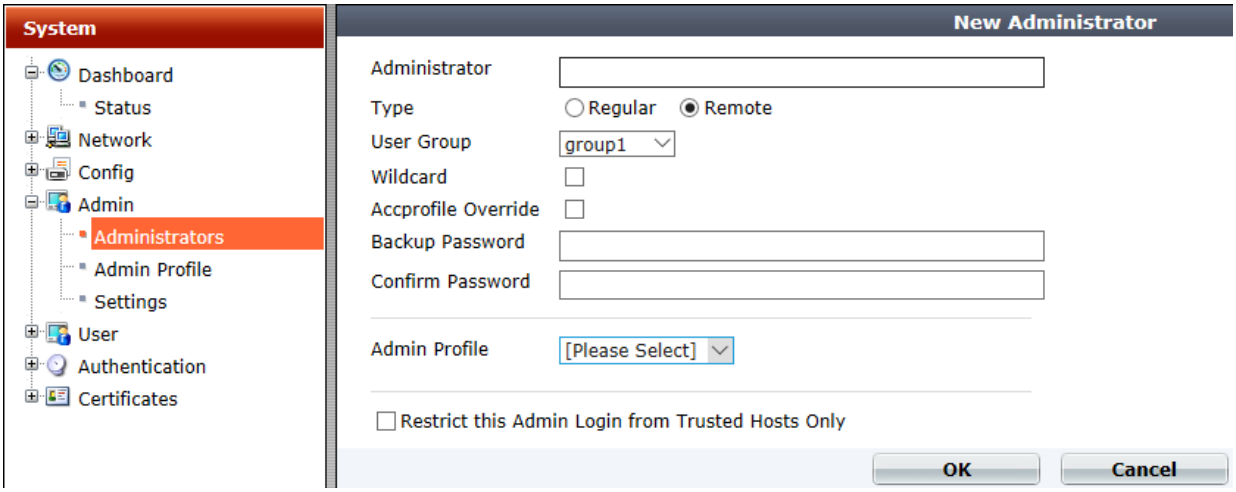

3. Enter the administrator name.
4. Select the type of account. If you select **Remote**, the system can reference a RADIUS or TACAS+ server.
5. When selecting Remote or PKI accounts, select the User Group the account will access.
6. Enter the password for the user. Passwords can be up to 256 characters in length.
7. Click **OK**.

Using the CLI:

```
config system admin
   edit <admin_name>
      set password <password>
      set accprofile <profile_name>
   end
```

Monitoring administrators

You can view the administrators logged in using the **System Information** widget on the **Dashboard**. On the widget is the **Current Administrator** row that shows the administrator logged in and the total logged in. Selecting **Details** displays the information for each administrator: where they are logging in from and how (CLI, Web-based manager) and when they logged in.

You are also able to monitor the activities the administrators perform using Event Logging. Event logs include a number of options to track configuration changes.

To set logging using the Web-based manager:

1. Go to **Log** > **Log Config** > **Log Setting**.

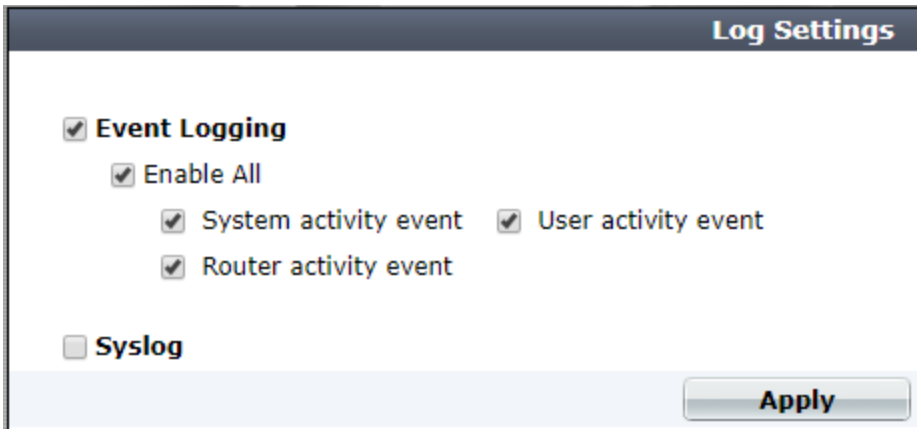

2. Under **Event Logging**, ensure that **System activity event** is selected.
3. Click **Apply**.

To set logging using the CLI:

```
config log eventfilter
   set event enable
   set system enable
end
```

To view the logs, go to **Log** > **Event Log** > **System**.

Setting the default administrator password

By default, your system has an administrator account set up with the user name `admin` and no password. To prevent unauthorized access, it is highly recommended that you add a password to this account.

To change the default password:

1. Go to **System** > **Administrators**.
2. Edit the **admin** account.
3. Select **Change Password**.

4. Leave **Old Password** blank, enter the **New Password** and re-enter the password for confirmation.
5. Select **OK**.

Setting the password retries and lockout time

By default, the system includes a set number of three password retries, allowing the administrator a maximum of three attempts to log into their account before they are locked out for a set amount of time (by default, 60 seconds).

The number of attempts can be set to an alternate value, as well as the default wait time before the administrator can try to enter a password again. You can also change this value to make it more difficult to hack. Both settings are must be configured with the CLI

To configure the lockout options:

```
config system global
   set admin-lockout-threshold <failed_attempts>
   set admin-lockout-duration <seconds>
end
```

For example, to set the lockout threshold to one attempt and the duration before the administrator can try again to log in to five minutes, enter these commands:

```
config system global
   set admin-lockout-threshold 1
   set admin-lockout-duration 300
end
```

Setting the idle timeout

By default, the GUI disconnects administrative sessions if no activity occurs for five minutes. This prevents someone from using the GUI if the management PC is left unattended.

To change the idle timeout

1. Go to **System** > **Admin** > **Settings**
2. Enter the time in minutes in the **Idle Timeout** field.
3. Update other settings as required:
   - TCP/UDP port values for HTTP, HTTPS, Telnet, SSH
   - Display language
4. Select **Apply**.

Configuring administrative logins

You can configure the RADIUS server to set the access profile. This process uses RADIUS vendor-specific attributes (VSAs) passed to the FortiSwitch for authorization. The RADIUS access profile override is mainly used for administrative logins.

Using the Web-based manager:

1. Go to **System > Admin > Administrators**.
2. Click **Remote**.

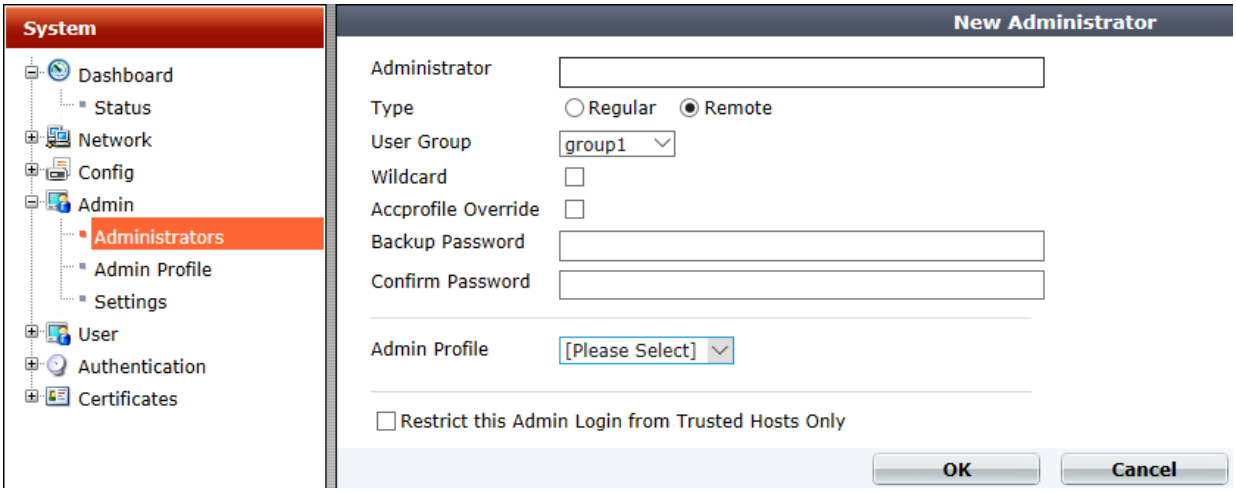

3. In the Administrator field, enter a name for the RADIUS-system administrator group.
4. Select the user group.
5. Click **Wildcard**.
6. Click **Accprofile Override**.
7. Click **OK**.

Using the CLI:

The following code creates a RADIUS-system admin group with accprofile-override enabled:

```
config system admin
   edit "RADIUS_Admins"
      set remote-auth enable
      set accprofile no_access
      set wildcard enable
      set remote-group "RADIUS_Admins"
      set accprofile-override enable
   next
```

Ensure that the RADIUS server is configured to send the appropriate VSA.

To send an appropriate group membership and access profile, set VSA 1 and VSA 6, as in the following code:

```
VENDOR fortinet 12356
ATTRIBUTE Fortinet-Group-Name 1 <admin profile>
ATTRIBUTE Fortinet-Access-Profile 6 <access profile>
```

The value of VSA 1 must match the remote group, and VSA 6 must match a valid access profile.

Configuring security checks

You can enable various security checks for incoming TCP/UDP packets. The packet is dropped if the system detects the specified condition. Use the appropriate syntax for your FortiSwitch model:

- Syntax (for models FS108D-POE and FS112D-POE) on page 42
- Syntax (for all other FortiSwitch models) on page 42

Syntax (for models FS108D-POE and FS112D-POE)

```
config switch security-feature
   set tcp-syn-data {enable | disable}
   set tcp-udp-port-zero {enable | disable}
   set tcp_flag_zero {enable | disable}
   set tcp_flag_FUP {enable | disable}
   set tcp_flag_SF {enable | disable}
   set tcp_flag_SR {enable | disable}
   set tcp_frag_ipv4_icmp {enable | disable}
   set tcp_arp_mac_mismatch {enable | disable}
```

Variable	Description	Default
tcp-syn-data	TCP SYN packet contains additional data (possible DoS attack).	disable
tcp-udp-port-zero	TCP or UDP packet has source or destination port set to zero.	disable
tcp_flag_zero	TCP packet with all flags set to zero.	disable
tcp_flag_FUP	TCP packet with FIN, URG and PSH flag set.	disable
tcp_flag_SF	TCP packet with SYN and FIN flag set.	disable
tcp_flag_SR	TCP packet with SYN and RST flag set.	disable
tcp_frag_ipv4_icmp	Fragmented ICMPv4 packet.	disable
tcp_arp_mac_mismatch	ARP packet with MAC source address mismatch between the layer- 2 header and the ARP packet payload.	disable

Syntax (for all other FortiSwitch models)

```
config switch security-feature
   set sip-eq-dip {enable | disable}
   set tcp-flag {enable | disable}
   set tcp-port-eq {enable | disable}
   set tcp-flag-FUP {enable | disable}
   set tcp-flag-SF {enable | disable}
   set v4-first-frag {enable | disable}
   set udp-port-eq {enable | disable}
   set tcp-hdr-partial {enable | disable}
   set macsa-eq-macda {enable | disable}
```

Variable	Description	Default
sip-eq-dip	TCP packet with source IP equal to destination IP.	disable
tcp_flag	DoS attack checking for TCP flags.	disable
tcp-port-eq	TCP packet with source and destination TCP port equal.	disable
tcp-flag-FUP	TCP packet with FIN, URG, and PSH flags set, and sequence number is zero.	disable
tcp-flag-SF	TCP packet with SYN and FIN flag set.	disable
v4-first-frag	DoS attack checking for IPv4 first fragment.	disable
udp-port-eq	IP packet with source and destination UDP port equal.	disable
tcp-hdr-partial	TCP packet with partial header.	disable
macsa-eq-macda	Packet with source MAC equal to destination MAC.	disable

Configuring SNMP

Simple Network Management Protocol (SNMP) enables you to monitor hardware on your network.

The FortiSwitch SNMP implementation is read-only. SNMP v1-compliant and v2c-compliant SNMP managers have read-only access to FortiSwitch system information through queries and can receive trap messages from the FortiSwitch.

To monitor FortiSwitch system information and receive FortiSwitch traps, you must first compile the Fortinet and FortiSwitch management information base (MIB) files. A MIB is a text file that describes a list of SNMP data objects that are used by the SNMP manager. These MIBs provide information that the SNMP manager needs to interpret the SNMP trap, event, and query messages sent by the FortiSwitch SNMP agent.

FortiSwitch core MIB files are available for download by going to **System** > **Config** > **SNMP** and selecting the MIB download link.

This chapter covers the following topics:

SNMP access

Ensure that the management VLAN has SNMP added to the access-profiles.

Using the Web-based manager:

1. Go to **System** > **Network** > **Interface**.
2. Edit the management interface.
3. Set **SNMP** in the access profiles.
4. Select **Apply**.

Using the CLI:

```
config system interface
   edit <name>
      set allowaccess <access_types>
   end
end
```

NOTE: Re-enter the existing allowed access types and add `snmp` to the list.

SNMP agent

Create the SNMP agent.

Using the Web-based manager:

1. Go to **System** > **Config** >**SNMP**.
2. Click **Enable** for the SNMP Agent.
3. Enter a descriptive name for the agent.
4. Enter the location of the FortiGate unit.
5. Enter a contact or administrator for the SNMP Agent or FortiSwitch unit.
6. Select **Apply**.

Using the CLI:

```
config system snmp sysinfo
   set status enable
   set contact-info <contact_information>
   set description <description_of_FortiSwitch>
   set location <FortiSwitch_location>
end
```

SNMP community

An SNMP community is a grouping of devices for network administration purposes. Within that SNMP community, devices can communicate by sending and receiving traps and other information. One device can belong to multiple communities, such as one administrator terminal monitoring both a FortiGate SNMP and a FortiSwitch SNMP community.

Add SNMP communities to your FortiSwitch so that SNMP managers can connect to view system information and receive SNMP traps.

You can add up to three SNMP communities. Each community can have a different configuration for SNMP queries and traps. Each community can be configured to monitor the FortiSwitch for a different set of events. You can also add the IP addresses of up to eight SNMP managers for each community.

Adding an SNMP v1/v2c community

Using the Web-based manager:

1. Go to **System** > **Config** > **SNMP**.
2. In the SNMP v1/v2c area, select **Create New**.
3. Enter a community name.
4. Enter the IP address and identify the SNMP managers that can use the settings in this SNMP community to monitor the FortiSwitch.
5. Select the interface if the SNMP manager is not on the same subnet as the FortiSwitch.
6. Enter the port number that the SNMP managers in this community use for SNMP v1 and SNMP v2c queries to receive configuration information from the FortiSwitch. Select the **Enable** check box to activate queries for each SNMP version.
7. Enter the local and remote port numbers that the FortiSwitch uses to send SNMP v1 and SNMP v2c traps to the SNMP managers in this community.

8. Select the **Enable** check box to activate traps for each SNMP version.
9. Select **OK**.

Using the CLI:

```
config system snmp community
   edit <index_number>
      set events <events_list>
      set name <community_name>
      set query-v1-port <port_number>
      set query-v1-status {enable | disable}
      set query-v2c-port <port_number>
      set query-v2c-status {enable | disable}
      set status {enable | disable}
      set trap-v1-lport <port_number>
      set trap-v1-rport <port_number>
      set trap-v1-status {enable | disable}
      set trap-v2c-lport <port_number>
      set trap-v2c-rport <port_number>
      set trap-v2c-status {enable | disable}
```

Adding an SNMP v3 community

Using the Web-based manager:

1. Go to **System** > **Config** > **SNMP**.
2. In the SNMP v3 area, select **Create New**.
3. Enter a User Name.
4. Select a Security Level and associated authorization algorithms.
5. Enter the Port number that the SNMP managers in this community use to receive configuration information from the FortiGate unit. Select the **Enable Query** check box to activate queries for each SNMP version.
6. Select the events to report.
7. Select **OK**.

Using the CLI:

```
config system snmp user
   edit <index_number>
      set events <event_selections>
      set queries enable
      set query-port <port_number>
      set security-level [auth-priv | auth-no-priv | no-auth-no-priv}
   end
```

Global system settings

This chapter covers the following topics:

Configuration file settings

You can set preferences for the configuration files:

1. Go to **System** > **Config** > **Settings**
2. Select a value for Configuration Save:
   - **Auto** - system automatically saves the configuration after each change.
   - **Manual** - you must manually save configuration changes, from **System** > **Config** > **Revisions**.
   - **Revertive** - you must manually save configuration changes. The system reverts to the saved configuration after a timeout. You can set the timeout using the CLI:
     ```
     config system global
     set cfg-revert-timeout <integer>
     ```
3. If you select **Revision Backup on Logout**, the FortiSwitch will create a configuration file each time a user logs out.
4. If you select **Revision Backup on Upgrade**, the FortiSwitch will create a configuration file before starting a system upgrade.
5. If you select **Strong Crypto**, the configuration is stored encrypted with strong cryptography.
6. Click **Apply**.

Configuration file revisions

Using the Web-based manager:

1. Go to **System** > **Config** > **Revisions**
   The system displays a new page with an entry for each configuration file revision.
2. When you select a revision, the following commands are available:
   - **Delete** - deletes the revision file.
   - **Details** - displays the contents of the revision file.
   - **Change** Comments - to edit the comments field for this revision file.

- **Revert** - reverts the system configuration to use this revision file.
- **Upload** - uploads the revision file to your local machine.

3. If you select two revision files, click **Diff** to display the differences between the two files.

Using the CLI:

Use the following command to display the list of configuration file revisions:

```
execute revision list config
```

The FortiSwitch assigns a numerical ID to each configuration file. To display a particular configuration file contents, use the following command and specify the ID of the configuration file:

```
execute revision show config id <ID number>
```

The following example displays the list of configuration file revisions:

```
# execute revision list config

ID TIME ADMIN FIRMWARE VERSION COMMENT
1 2015-08-31 11:11:00 admin V3.0.0-build117-REL0 Automatic backup (session
expired)
2 1969-12-31 16:06:29 admin V3.0.0-build150-REL0 baseline
3 2015-08-31 15:19:31 admin V3.0.0-build150-REL0 baseline
4 2015-08-31 15:28:00 admin V3.0.0-build150-REL0 with admin timeout
```

The following example displays the configuration file contents for revision ID 62:

```
# execute revision show config id 62

#config-version=FS1D24-3.04-FW-build171-160201:opmode=0:vdom=0:user=admin
#conf_file_ver=1784779075679102577
#buildno=0171
#global_vdom=1
config system global
   set admin-concurrent enable
       ...
   (output truncated)
```

IP conflict detection

IP conflicts can occur when two systems on the same network are using the same IP address. FortiSwitch monitors the network for conflicts and raises a system log message and an SNMP trap when it detects a conflict.

The IP conflict detection feature provides two methods to detect a conflict. The first method relies on a remote device to send a broadcast ARP (Address Resolution Protocol) packet claiming ownership of a particular IP address. If the IP address in the source field of that ARP packet matches any of the system interfaces associated with the receiving FortiSwitch system, the system logs a message and raises an SNMP trap.

For the second method, the FortiSwitch actively broadcasts gratuitous ARP packets when any of the following events occurs:

- System boot-up
- Interface status changes from down to up
- IP address change

If a system is using the same IP address, the FortiSwitch will receive a reply to the gratuitous ARP. If it receives a reply, the system logs a message.

Configuring IP conflict detection

IP conflict detection is enabled on a global basis. The default setting is enabled.

Using the Web-based manager:

1. Go to **Network > Settings**.
2. Set **IP Conflict Detection**
3. Select **OK.**

Using the CLI:

```
config system global
   set detect-ip-conflict <enable|disable>
```

Viewing IP conflict detection

If the system detects an IP conflict, the system generates the following log message:

```
IP Conflict: conflict detected on system interface mgmt for IP address 10.10.10.1
```

Port flap guard

A flapping port can create instability in protocols such as STP. If a port is flapping, STP must continually recalculate the role for each port.

The port flap guard feature will detect a flapping port, and the system will shut down the port if necessary. You can manually reset the port and restore it to the enabled state.

Configuring port flap guard

Port flap guard is configured and enabled on a global basis. The default setting is disabled. Flap rate ranges from 5 to 300.

Using the Web-based manager:

1. Go to **Switch > Flap Guard > Settings**.
2. Enable **Flap Guard**.
3. Enter a value for **Flap duration** and **Flap rate**.
4. Click **Apply** to save the changes.

Using the CLI:

```
config switch flapguard settings
   set status [ disable | enable ]
   set flap-rate <integer>
   set flap-duration <integer>
```

Use the following command to reset a port and restore it to service:

```
execute flapguard reset <port>
```

Viewing the port flap guard configuration

Display the status of the port flap guard configuration using the following command:

```
show switch flapguard settings
```

Display the port flap guard information for each port using the following command:

```
diagnose flapguard instance status
```

Link monitor

You can monitor the link to a server. The FortiSwitch sends periodic ping messages to test that the server is available.

Configuring the link monitor

Using the Web-based manager:

1. Go to **Router > Link Monitor > Probes**.
2. Click **Create New** to create a new probe.
3. Enter an IP address for the **Gateway IP**.
4. Configure the other fields as required (see the table in this section for field descriptions).
5. Click **Advance Settings** to view additional fields that you can configure.
6. Click **OK** to save the changes.

Using the CLI:

```
config system link-monitor
   edit "1"
      set srcintf <string>
      set protocol (arp | ping)
      set gateway-ip <IP address>
      set source-ip <IP address>
      set interval <integer>
      set timeout <integer>
      set failtime <integer>
      set recoverytime <integer>
      set update-cascade-interface (enable | disable)
      set update-static-route (enable | disable)
```

```
      set status (enable | disable)
   next
end
```

Variable	Description
srcintf	Interface where the monitor traffic is sent.
protocol	Protocols used to detect the server. Select ARP or ping.
gateway-ip	Gateway IP used to PING the server.
source-ip	Source IP used in packet to the server.
interval	Detection interval in seconds. The range is 1-3600.
timeout	Detect request timeout in seconds. The range is 1-255.
failtime	Number of retry attempts before bringing the server down. The range is 1-10.
recoverytime	Number of retry attempts before bringing the server up. The range is 1-10.
update-cascade-interface	Enable or disable update cascade interface.
update-static-route	Enable or disable update static route.
status	Enable or disable link monitor administrative status.

Unicast hashing

You can configure the trunk hashing algorithm for unicast packets to use the source port:

```
config switch global
   set trunk-hash-unicast-src-port {enable | disable}
end
```

Cut-through switching mode

By default, all FortiSwitch models use the store-and-forward technique to forward packets. This technique waits until the entire packet is received, verifies the content, and then forwards the packet.

The FSW-1024D, FSW-1048D, and FSW-3032D models also have a cut-through switching mode to reduce latency. This technique forwards the packet as soon as the switch receives it.

NOTE: For the FSW-3032D model, the cut-through switching mode is not supported on split ports.

To change the switching mode for the main buffer for these three models, use the following commands:

```
config switch global
```

```
   set packet-buffer-mode {store-forward | cut-through}
end
```

NOTE: Changing the switching mode might stop traffic on all ports during the change.

Enabling packet forwarding

NOTE: These commands apply only to the 124D, 124D-POE, 200 Series, and 400 Series.

If you want to use layer-3 interfaces and IGMP snooping on certain FortiSwitch models, you must enable the forwarding of reserved multicast packets and IPv6 neighbor-discovery packets to the CPU. These features are enabled by default.

```
config switch global
   set reserved-mcast-to-cpu {enable | disable}
   set neighbor-discovery-to-cpu {enable | disable}
end
```

Physical port settings

This following sections describe the configuration settings that are associated with FortiSwitch physical ports:

Configuring general port settings

Using the Web-based manager:

1. Go to **Switch > Physical > Interface** and select the port to update.
2. Enter values for Name and Description.
3. Select the Admin port status.
4. Select **OK.**

Using the CLI:

```
config switch physical-port
   edit <port>
      set description <string>
      set max-frame-size
      set status (up | down)
```

General port settings include:

- **description** - Text description for the port
- **max-frame-size** - Maximum frame size in bytes (between 68 and 9216)
- **status** - Administrative status of the port

Viewing port statistics

Use the following command to get statistics for a specific port:

```
diag switch physical-ports port-stats <port_number>
```

For example:

```
diag switch physical-ports port-stats port1
```

Configuring flow control

Flow control represents the ability to configure a port to send or receive a “pause frame” (that is, a special packet that signals a source to stop sending flows for a specific time interval because the buffer is full). By default, flow control is disabled on all ports.

```
config switch physical-port
   edit <port>
      set flow-control (both | rx | tx | disable)
```

Parameters enable flow control to do the following:

- **rx** - receive pause control frames
- **tx** - transmit pause control frames
- **both** - transmit and receive pause control frames

Auto-module speed detection

When you enable auto-module speed detection, the system reads information from the module and sets the port speed to the maximum speed that is advertised by the module. If the system encounters a problem when reading from the module, it sets the default speed (default value is platform specific).

When auto-module sets the speed, the system creates a log entry noting this speed.

NOTE: Auto-speed detection is supported on 1/10G ports, but not on higher speed ports (such as 40G).

Setting port speed (autonegotiation)

By default, all of the FortiSwitch user ports are set to autonegotiate the port speed. You can also manually set the port speed:

Using the Web-based manager:

1. Go to **Switch > Port > Physical** and select the port.
2. Click **Edit**.
3. Select the desired port speed.
4. Click **OK**.

Using the CLI:

```
config switch physical-port
   edit <port>
      set speed (auto | 10full | 10half | 100full | 100half | 1000auto)
   end
end
```

Viewing auto-module configuration

Display the status of auto-module using following command:

```
config switch physical-port
   edit port47
      show
   config switch physical-port
      edit "port47"
         set max-frame-size 16360
         set speed 10000full
      next
   end
get
   name : port47
   description : (null)
   flow-control : both
   link-status : down
   lldp-transmit : disable
   max-frame-size : 16360
   port-index : 47
   speed : 10000full
   status : up
```

Link-layer discovery protocol

The Fortinet data center switches support LLDP (transmission and reception). The link layer discovery protocol (LLDP) is a vendor-neutral layer-2 protocol that enables devices on a layer-2 segment to discover information about each other.

For details, refer to LLDP-MED on page 89.

Configuring power over Ethernet

Power over Ethernet (PoE) describes any system that passes electric power along with data on twisted pair Ethernet cabling. Doing this allows a single cable to provide both data connection and electric power to devices (for example, wireless access points, IP cameras, and VoIP phones).

PoE is only available on models with the POE suffix in the model number (for example, FS-108D-POE).

Enabling PoE on a port

```
config switch physical-port
   edit <port>
      set poe-status enable
      set poe-pre-standard-detection {enable | disable}
```

```
      set poe-reset reset
   end
end
```

Determining the PoE power capacity

To determine the PoE power capacity, use the following command:

```
get switch poe inline
```

Reset the PoE power on a port

To reset the PoE power on a port, use the following command:

```
execute poe-reset <port>
```

Selecting how power is allocated

When power to PoE ports is allocated by priority, lower numbered ports have higher priority so that port 1 has the highest priority. When more power is needed than is available, higher numbered ports are disabled first.

When power to PoE ports is allocated by first-come, first-served (FCFS), connected PoE devices receive power, but new devices do not receive power if there is not enough power.

If both priority power allocation and FCFS power allocation are selected, the physical port setting takes precedence over the global setting.

To select priority power allocation on a global basis, use the following command:

```
config switch global
   set poe-port-mode priority
end
```

To select FCFS power allocation on a global basis, use the following command:

```
config switch global
   set poe-port-mode first-come-first-served
end
```

To set the priority (from low to critical) for priority power allocation for a specific port, use the following command:

```
config switch physical-port
   edit <port>
      set poe-port-priority <priority>
```

Configure PoE with dynamic guard band (DGB)

The dynamic guard band is set automatically to the expected power of a port before turning on the port. So, when a PoE device is plugged in, the dynamic guard band is set to the maximum power of the device type based on the AF or AT mode. The AF mode DGB is 15.4 W, and the AT mode DGB is 36 W. When the FortiSwitch is fully loaded, the dynamic guard band prevents a new PoE device from turning on.

To avoid this issue, change the port mode using the following commands:

```
config switch physical-port
   edit <port>
```

```
    set port-mode IEEE802_3AF
end
```

Display PoE information for a port

To display PoE information for a port, use the following command:

```
diagnose switch poe status <port>
```

The following example displays the information for port 6:

```
diagnose switch poe status port6
Port(6) Power:4.20W, Power-Status: Delivering Power
Power-Up Mode: Normal Mode
Remote Power Device Type: IEEE802.3AT PD
Power Class: 4
Defined Max Power: 30.0W, Priority:3
Voltage: 54.00V
Current: 71mA
```

Diagnostic monitoring interface module status

With diagnostic monitoring interface (DMI), you can view the following information

- Module details (detail)
- Eeprom contents (eeprom)
- Module limits (limit)
- Module status (status)
- Summary information of all a port's modules (summary)

DMI is supported on all models except FortiSwitch 124D.

Using the Web-based manager:

Go to **Switch > Monitor > Modules**.

Using the CLI:

Use the following commands to enable or disable DMI status for the port. If you set the status to `global`, the port setting will match the global setting:

```
config switch physical-port
   edit <interface>
      set dmi-status {disable | enable | global}
```

Use the `get switch modules detail/status` command to display DMI information:

```
FS108D3W14000720 # get switch modules detail port10
```

```
Port(port10)
identifier SFP/SFP+
connector Unk (0x00)
transceiver 1000-Base-T
encoding 8B/10B
Length Decode Common
length_smf_1km N/A
length_cable 100 meter
SFP Specific
length_smf_100m N/A
length_50um_om2 N/A
length_62um_om1 N/A
length_50um_om3 N/A
vendor FINISAR CORP.
vendor_oid 0x009065
vendor_pn FCLF-8521-3
vendor_rev A
vendor_sn PBR1X35
manuf_date 06/20/2007
```

The following is an example of the output for the `switch modules status` command:

```
FS108D3W14000720 # get switch modules status port9

_____________________________________________________________
Port(port9)
alarm_flags 0x0040
warning_flags 0x0040
temperature 18.792969 C
voltage 3.315100 volts
laser_bias 0.750800 mAmps
tx_power -2.502637 dBm
rx_power -40.000000 dBm
options 0x000F ( TX_DISABLE TX_FAULT RX_LOSS TX_POWER_LEVEL1 )
options_status 0x000C ( RX_LOSS TX_POWER_LEVEL1 )
```

Configuring split port

On FortiSwitch models that provide 40G QSFP (quad small form-factor pluggable) interfaces, you can install a breakout cable to convert one 40G interface into four 10G interfaces.

Notes

- Split port is supported on the following FortiSwitch models:
  - 3032D (port5 to port28 are splittable)
  - 524D, 524D-FPOE (port29 and port30 are splittable)
  - 548D, 548D-FPOE (port53 and port54 are splittable)
- Currently, the maximum number of ports supported in software is 64. Therefore, only 10 QSFP ports can be split. This limitation applies to all of the models, but only the 3032D has enough ports to encounter this limit.
- Split port is not supported in FortiLink mode (that is, FortiSwitch managed by FortiGate).

Configuring split port

Use the following commands to configure split port:

```
config switch phy-mode
   set port-configuration <default | disable-port54 | disable-port41-48>
   set <port-name>-phy-mode <1x40G | 4x10G>
   ...
   (one entry for each port that supports split port)
end
```

NOTE: The `port-configuration` command applies solely to the 548D and 548D-FPOE models.

The following settings are available:

- **disable-port54** - only port53 is splittable; port54 is unavailable.
- **disable-port41-48** - port41 to port48 are unavailable, but you can configure port53 and port54 in split-mode.

In the following example, a FortiSwitch 3032D is configured with ports 10, 14, and 28 set to 4x10G:

```
config switch phy-mode
   set port5-phy-mode 1x40G
   set port6-phy-mode 1x40G
   set port7-phy-mode 1x40G
   set port8-phy-mode 1x40G
   set port9-phy-mode 1x40G
   set port10-phy-mode 4x10G
   set port11-phy-mode 1x40G
   set port12-phy-mode 1x40G
   set port13-phy-mode 1x40G
   set port14-phy-mode 4x10G
   set port15-phy-mode 1x40G
   set port16-phy-mode 1x40G
   set port17-phy-mode 1x40G
   set port18-phy-mode 1x40G
   set port19-phy-mode 1x40G
   set port20-phy-mode 1x40G
   set port21-phy-mode 1x40G
   set port22-phy-mode 1x40G
   set port23-phy-mode 1x40G
   set port24-phy-mode 1x40G
   set port25-phy-mode 1x40G
   set port26-phy-mode 1x40G
   set port27-phy-mode 1x40G
   set port28-phy-mode 4x10G
end
```

The system applies the configuration only after you enter the `end` command, displaying the following message:

```
This change will cause a ports to be added and removed, this will cause loss of
    configuration on removed ports. The system will have to reboot to apply this change.
Do you want to continue? (y/n)y
```

To configure one of the split ports, use the notation ".x" to specify the split port:

```
config switch physical-port
   edit "port1"
```

```
      set lldp-profile "default-auto-isl"
      set speed 40000full
   next
   edit "port2"
      set lldp-profile "default-auto-isl"
      set speed 40000full
   next
   edit "port3"
      set lldp-profile "default-auto-isl"
      set speed 40000full
   next
   edit "port4"
      set lldp-profile "default-auto-isl"
      set speed 40000full
   next
   edit "port5.1"
      set speed 10000full
   next
   edit "port5.2"
      set speed 10000full
   next
   edit "port5.3"
      set speed 10000full
   next
   edit "port5.4"
      set speed 10000full
   next
```

Configuring QSFP low-power mode

On FortiSwitch models with QSFP (quad small form-factor pluggable) ports, you can enable or disable the low-power mode with the following CLI commands:

```
config switch physical-port
   edit <port_name>
      set qsfp-low-power-mode {enabled | disabled}
   end
end
```

For example:

```
config switch physical-port
   edit port12
      set qsfp-low-power-mode disabled
   end
end
```

Layer-2 interfaces

This chapter covers the following topics:

Switched interfaces

Default configuration will suffice for regular switch ports. By default, VLAN is set to 1, STP is enabled, and all other optional capabilities are disabled.

You can configure optional capabilities such as Spanning Tree Protocol, sFlow , 802.1x authentication, and Private VLANs. These capabilities are covered in subsequent sections of this document.

Using the Web-based manager:

1. Go to **Switch > Interface > Interface.**
2. Select the port to update and click **Edit**.
3. Select one or more ports to update and click **Edit**.
4. If you selected more than one port, the port names are displayed in the name field, separated by commas.
5. Enter new values as required for Native VLAN, Allowed VLANs and Untagged VLANs.
6. Click **OK** to save the changes.

Using the CLI:

```
config switch interface
   edit <port>
      set native-vlan <vlan>
      set allowed-vlans <vlan> [<vlan>] [<vlan> - <vlan>]
      set untagged-vlans <vlan> [<vlan>] [<vlan> - <vlan>]
      set stp-state {enabled | disabled}
      set edge-port {enabled | disabled}
      set security-mode {none| dot1x}
```

Viewing interface configuration

Display port configuration using the following command:

```
show switch interface <port>
```

Display port settings using following command:

```
config switch interface
   edit <port>
```

```
get
```

Dynamic MAC address learning

You can enable or disable dynamic MAC address learning on a port. The existing dynamic MAC entries are flushed when you change this setting. If you disable MAC address learning, you can set the behavior for an incoming packet with an unknown MAC address (to drop or forward the packet).

You can limit the number of learned MAC addresses on an interface or VLAN. The limit ranges from 1 to 128. If the learning limit is set to zero (the default), no limit exists. When the limit is exceeded, FortiSwitch adds a warning to the system log.

NOTE: Static MAC addresses are not counted in the limit. The limit refers only to learned MAC addresses.

Use the following CLI commands to configure dynamic MAC address learning:

```
config switch physical-port
   edit <port>
      set l2-learning (enable | disable)
      set l2-unknown (drop | forward)
   end
config switch interface
   edit <port>
      set learning-limit <0-128>
   end
config switch vlan
   edit <VLAN_ID>
      set learning-limit <0-128>
   end
```

NOTE: If you enable 802.1x MAC-based authorization on a port, you cannot change the `l2-learning` setting.

By default, each learned MAC address is aged out after 300 seconds. The value ranges from 10 to 1000,000 seconds. Set the value to zero to disable MAC address aging.

Use the following command to change this value:

```
config switch global
   set mac-aging-interval 200
end
```

If you want to see the first MAC address that exceeded a learning limit for an interface or VLAN, you can enable the learning limit violation log for a FortiSwitch. Only one violation is recorded per interface or VLAN.

To enable or disable the learning limit violation log, use the following commands. By default, the learning limit violation log is disabled.

```
config switch global
   set log-mac-limit-violations {enable | disable}
end
```

To view the content of the learning limit violation log, use one of the following commands:

- `get switch mac-limit-violations all`—to see the first MAC address that exceeded the learning limit on any interface or VLAN. An asterisk by the interface name indicates that the interface-based learning limit was

exceeded. An asterisk by the VLAN identifier indicates the VLAN-based learning limit was exceeded.

- `get switch mac-limit-violations interface <interface_name>`—to see the first MAC address that exceeded the learning limit on a specific interface
- `get switch mac-limit-violations vlan <VLAN_ID>`—to see the first MAC address that exceeded the learning limit on a specific VLAN

To reset the learning limit violation log, use one of the following commands:

- `execute mac-limit-violation reset all`—to clear all learning limit violation logs
- `execute mac-limit-violation reset interface <interface_name>`—to clear the learning limit violation log for a specific interface
- `execute mac-limit-violation reset vlan <VLAN_ID>`—to clear the learning limit violation log for a specific VLAN

Persistent (sticky) MAC addresses

You can make dynamically learned MAC addresses persistent when the status of a FortiSwitch port changes (goes down or up). By default, MAC addresses are not persistent.

NOTE: You cannot use persistent MAC addresses with 802.1x authentication.

Use the following command to configure the persistence of MAC addresses on an interface:

```
config switch interface
   edit <port>
      set sticky-mac <enable | disable>
   next
end
```

You can also save persistent MAC addresses to the FortiSwitch configuration file so that they are automatically loaded when the FortiSwitch is rebooted. By default, persistent entries are lost when a FortiSwitch is rebooted. Use the following command to save persistent MAC addresses for a specific interface or all interfaces:

```
execute sticky-mac save {all | interface <interface_name>}
```

Use the following command to delete the persistent MAC addresses instead of saving them in the FortiSwitch configuration file:

```
execute sticky-mac delete-unsaved {all | interface <interface_name>}
```

Static MAC addresses

You can configure one or more static MAC addresses on an interface.

Using the Web-based manager:

1. Go to **Switch > Static L2 > Entries**.
2. Click **Create** to create a new item.
3. Select an interface and enter a value for **MAC Address** and **VLAN ID**.
4. Click **Apply** to save the changes.

Using the CLI:

```
config switch static-mac
   edit <sequence_number>
      set description <optional_string>
      set interface <interface_name>
      set mac <static_MAC_address>
      set type {sticky | static}
      set vlan-id <VLAN_ID>
   end
```

For example:

```
config switch static-mac
   edit 1
      set description "first static MAC address"
      set interface port10
      set mac d6:dd:25:be:2c:43
      set type static
      set vlan-id 10
   end
```

Fortinet loop guard

A loop in a layer-2 network results in broadcast storms that have far-reaching and unwanted effects. Fortinet loop guard helps to prevent loops. When loop guard is enabled on a switch port, the port monitors its subtending network for any downstream loops.

The loop guard feature is designed to work in concert with STP rather than as a replacement for STP. Each port that has loop guard enabled will periodically broadcast loop guard data packets (LGDP) packets to its network. If a broadcast packet is subsequently received by the sending port, a loop exists downstream.

NOTE: If a port detects a loop, the system takes the port out of service to protect the overall network. The port returns to service after a configured timeout duration. If the timeout value is zero, you must manually reset the port.

By default, loop guard is disabled on all ports, and the timeout is set to zero.

Configuring loop guard

Using the Web-based manager:

1. Go to **Switch > Interface > Interface** or **Switch > Interface > Trunk**.
2. Select the port to update and click **Edit**.
3. Select one or more ports to update and click **Edit**.
4. If you selected more than one port, the port names are displayed in the name field, separated by commas.
5. Click **Enable Loop Guard**.
6. Click **OK**.

Using the CLI:

```
config switch interface
   edit port <number>
```

```
set loop-guard <enabled|disabled>
set loop-guard-timeout <integer>
```

When loop guard takes a port out of service, the system creates the following log messages:

```
Loop Guard: loop detected on <port_name>. Shutting down <port_name>
```

Use the following command to reset a port that detected a loop:

```
execute loop-guard reset <port>
```

Viewing the loop guard configuration

Use the following command to display the loop guard status for all ports:

```
diagnose loop-guard instance status
```

VLANs and VLAN tagging

FortiSwitch ports will process tagged and untagged Ethernet frames. Untagged frames do not carry any VLAN information.

Tagged frames include an additional header (the 802.1Q header) after the Source MAC address. This header includes a VLAN ID. This allows the VLAN value to be transmitted between switches.

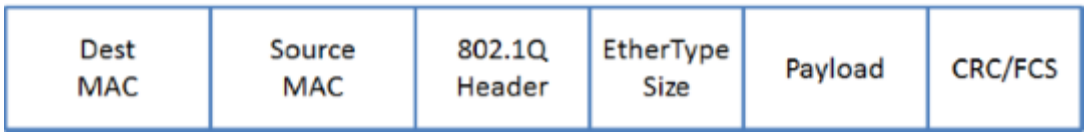

The FortiSwitch provides port parameters to configure and manage VLAN tagging.

This chapter covers the following topics:

Native VLAN

You can configure a native VLAN for each port. The native VLAN is like a default VLAN for untagged incoming packets. Outgoing packets for the native VLAN are sent as untagged frames.

The native VLAN is assigned to any untagged packet arriving at an ingress port.

At an egress port, if the packet tag matches the native VLAN, the packet is sent out without the VLAN header.

Allowed VLAN list

The allowed VLAN list for each port specifies the VLAN tag values for which the port can transmit or receive packets.

For a tagged packet arriving at an ingress port, the tag value must match a VLAN on the allowed VLAN list or the native VLAN.

At an egress port, the packet tag must match the native VLAN or a VLAN on the allowed VLAN list.

Untagged VLAN list

The untagged VLAN list on a port specifies the VLAN tag values for which the port will transmit packets without the VLAN tag. Any VLAN in the untagged VLAN list must also be a member of the allowed VLAN list.

The untagged VLAN list applies only to egress traffic on a port.

Packet processing

Ingress processing ensures that the port accepts only packets with allowed VLAN values (untagged packets are assigned the native VLAN, which is implicitly allowed). At this point, all packets are now tagged with a valid VLAN.

The packet is sent to each egress port that can send the packet (because the packet tag value matches the native VLAN or an Allowed VLAN on the port).

Ingress port

Untagged packet

- packet is tagged with the native VLAN and allowed to proceed
- the Allowed VLAN list is ignored

Tagged packet

- tag VLAN value must match an Allowed VLAN or the native VLAN
- packet retains the VLAN tag and is allowed to proceed

To control what types of frames are accepted by the port, use the following commands:

```
config switch interface
   edit <interface>
      set discard-mode <all-tagged | all-untagged | none>
   end
```

Variable	Description
all-tagged	Tagged frames are discarded, and untagged frames can enter the switch.
all-untagged	Untagged frames are discarded, and tagged frames can enter the switch.
none	By default, all frames can enter the switch, and no frames are discarded.

Egress port

All packets that arrive at an egress port are tagged packets.

If the packet tag value is on the Allowed VLAN list, the packet is sent out with the existing tag.

If the packet tag value is the native VLAN or on the Untagged VLAN list, the tag is stripped, and then the packet is sent out.

Otherwise, the packet is dropped.

Configuring VLANs

Use the following steps to create a new VLAN interface:

Using the Web-based manager:

1. Go to **System > Network > Interface** and select **Create New** to create a VLAN.
2. Give the VLAN an appropriate name.
3. Set **Interface** to **internal**.
4. Set a **VLAN ID**.
5. Assign an **IP/Netmask**.
6. Set **Administrative Access** to use the desired protocols to connect to the interface.
7. Select **OK**.

Using the CLI:

```
config system interface
   edit <vlan name>
      set ip <address>
      set allowaccess <access_types>
      set switch-members <port>
      set vlanid <VLAN id>
   end
end
```

Example 1

Example flows for tagged and untagged packets.

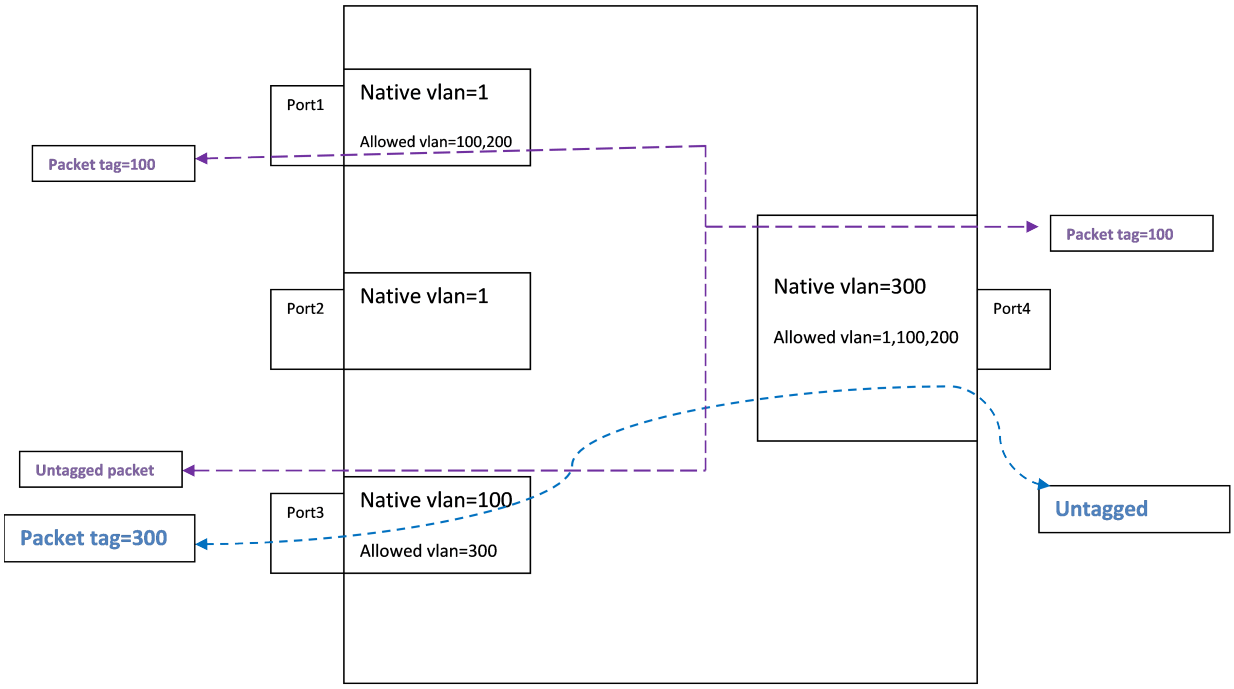

Purple flow

An untagged packet arriving at Port3 is assigned VLAN 100 (the native VLAN) and flows to all egress ports that will send VLAN 100 (Port1 and Port4).

A tagged packet (VLAN 100) arriving at Port4 is allowed (VLAN 100 is allowed). The packet is sent out from Port1 and Port3. On Port3, VLAN 100 is the native VLAN, so the packet is sent without a VLAN tag.

Blue flow

An untagged packet arriving at Port 4 is assigned VLAN 300 (the native VLAN). Then it flows out all ports that will send Vlan300 (Port 3).

A tagged packet (VLAN 300) arriving at Port3 is allowed. The packet is sent to egress from Port4. VLAN 300 is the native VLAN on Port4, so the packet is sent without a VLAN tag.

Example 2

Example of invalid tagged VLAN.

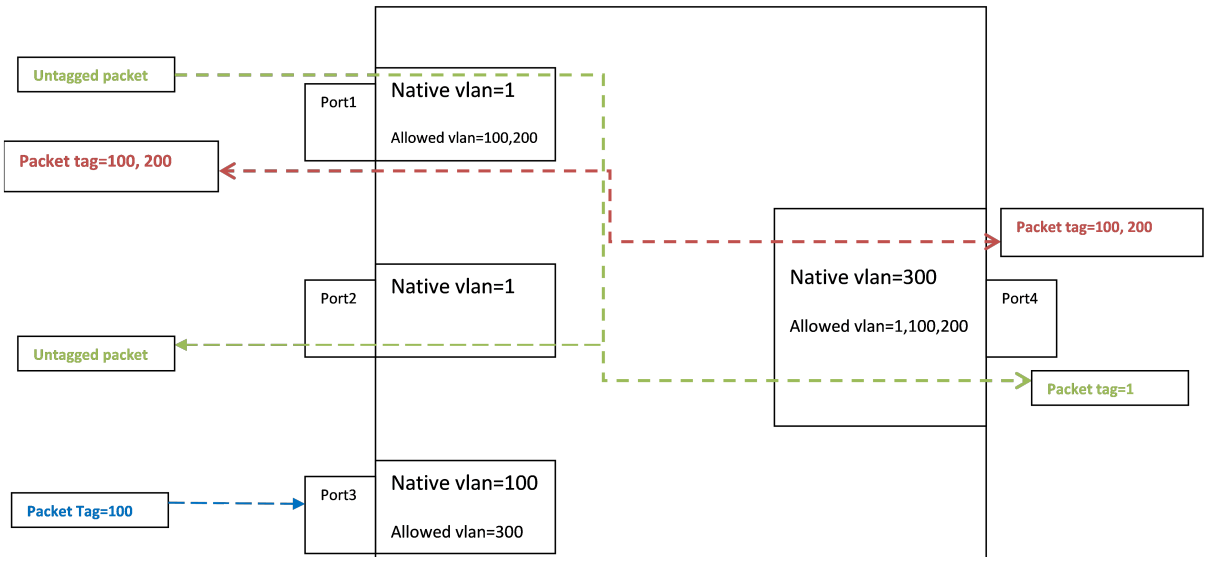

Green flow

Between Port1 and Port2, packets are assigned to VLAN 1 at ingress, and then the tag is removed at egress.

Blue flow

Incoming on Port 3, a tagged packet with VLAN value 100 is allowed, because 100 is the Port 3 native VLAN (the hardware VLAN table accepts a tagged or untagged match to a valid VLAN).

The packet will be sent on port1 and port4 (with packet tag 100).

Spanning Tree Protocol

FortiSwitch supports Spanning Tree Protocol (a link-management protocol that ensures a loop-free layer-2 network topology) as well as Multiple Spanning Tree Protocol (MSTP), which is defined in the IEEE 802.1Q standard.

This chapter covers the following topics:

MSTP overview and terminology

MSTP supports multiple spanning tree instances, where each instance carries traffic for one or more VLANs (the mapping of VLANs to instances is configurable).

MSTP is backward-compatible with STP and Rapid Spanning Tree Protocol (RSTP). A layer-2 network can contain switches that are running MSTP, STP, or RSTP.

MSTP is built on RSTP, so it provides fast recovery from network faults and fast convergence times.

Regions

A region is a set of interconnected switches that have the same multiple spanning tree (MST) configuration (region name, MST revision number, and VLAN-to-instance mapping). A network can have any number of regions. Regions are independent of each other because the VLAN-to-instance mapping is different in each region.

FortiSwitch supports 15 MST instances in a region. Multiple VLANs can be mapped to each MST instance. Each switch in the region must have the identical mapping of VLANs to instances.

The MST region acts like a single bridge to adjacent MST regions and to non-MST STPs.

IST

Instance 0 is a special instance, called the internal spanning-tree instance (IST). IST is a spanning tree that connects all of the MST switches in a region. All VLANs are assigned to the IST.

IST is the only instance that exchanges bridge protocol data units (BPDUs). The MSTP BPDU contains information for each MSTP instance (captured in an M-record). The M-records are added to the end of a regular RSTP BPDU. This allows MSTP region to inter-operate with an RSTP switch.

CST

The common spanning tree (CST) interconnects the MST regions and all instances of STP or RSTP that are running in the network.

Hop count and message age

MST does not use the BPDU message age within a region. The message-age and maximum-age fields in the BPDU are propagated unchanged within the region.

Within the region, a hop-count mechanism is used to age out the BPDU. The IST root sends out BPDUs with the hop count set to the maximum number of hops. The hop count is decremented each time the BPDU is forwarded. If the hop count reaches zero, the switch discards the BPDU and ages out the information on the receiving port.

STP port roles

STP assigns a port role to each switch port. The role is based on configuration, topology, relative position of the port in the topology, and other considerations. Based on the port role, the port either sends or receives STP BPDUs and forwards or blocks the data traffic. Here is a brief summary of each STP port role:

- **Designated**—One designated port is elected per link (segment). The designated port is the port closest to the root bridge. This port sends BPDUs on the link (segment) and forwards traffic towards the root bridge. In an STP converged network, each designated port is in the STP forwarding state.
- **Root**—The bridge can have only one root port. The root port is the port that leads to the root bridge. In an STP converged network, the root port is in the STP forwarding state.
- **Alternate**—Alternate ports lead to the root bridge but are not root ports. The alternate ports maintain the STP blocking state.
- **Backup**—This is a special case when two or more ports of the same switch are connected together (either directly or through shared media). In this case, one port is designated, and the remaining ports are backup (in the STP blocking state).

STP loop protection

The STP loop-protection feature provides additional protection against layer-2 forwarding loops (STP loops). An STP loop is created when an STP blocking port in a redundant topology erroneously transitions to the forwarding state.

A port remains in blocking state only if it continues to receive BPDU messages. If it stops receiving BPDUs (for example, due to unidirectional link failure), the blocking port (alternate or backup port) becomes designated and transitions to a forwarding state. In a redundant topology, this situation may create a loop.

If the loop-protection feature is enabled on a port, that port is forced to remain in blocking state, even if the port stops receiving BPDU messages. It will not transition to forwarding state and does not forward any user traffic.

The loop-protection feature is enabled on a per-port basis. Fortinet recommends that you enable loop protection on all nondesignated ports (all root, alternate, and backup ports).

STP root guard

Root guard protects the interface on which it is enabled from becoming the path to root. When enabled on an interface, superior BPDUs received on that interface are ignored or dropped. Without using root guard, any switch that participates in STP maintains the ability to reroute the path to root. Rerouting might cause your network to transmit large amounts of traffic across suboptimal links or allow a malicious or misconfigured device to pose a security risk by passing core traffic through an insecure device for packet capture or inspection. By enabling root guard on multiple interfaces, you can create a perimeter around your existing paths to root to enforce the specified network topology.

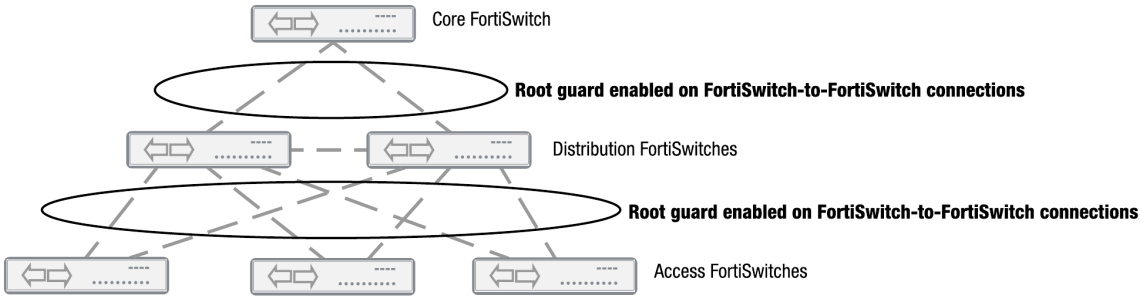

STP BPDU guard

Similar to root guard, BPDU guard protects the designed network topology. When BPDU guard is enabled on STP edge ports, any BPDUs received cause the ports to go down for a specified number of minutes. The BPDUs are not forwarded, and the network edge is enforced.

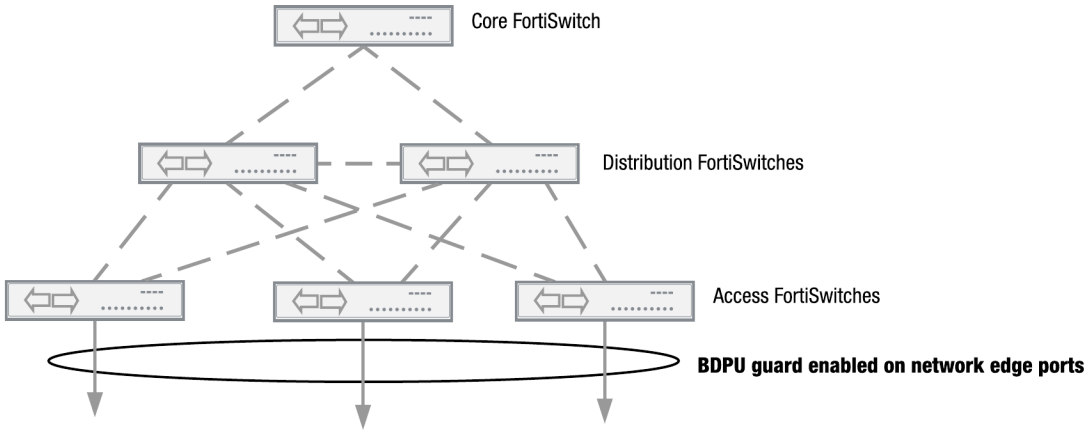

MSTP configuration

MSTP configuration consists of the following steps:

1. Configure STP settings that are common to all MST instances.
2. Configure settings that are specific to each MST instance.
3. Configure loop-protection on all nondesignated ports.

Configuring STP settings

Some STP settings (region name and MST revision number) are common to all MST instances. Also, protocol timers are common to all instances because only the IST sends out BPDUs.

Using the Web-based manager:

1. Go to **Switch** > **STP** > **Settings.**
2. Update the settings as described in the following table.
3. Click **Apply** to save the settings.

Settings	Guidelines
Enable	Enables MSTP for this switch.
Name	Region name. All switches in the MST region must have the identical name.
Revision	The MSTP revision number. All switches in the region must have the same revision number. Range of values is 0 - 65535. Default value is 0.
Hello-Time	Hello time is how often (in seconds) that the switch sends out a BPDU. Range of values is 1 to 10. Default value is 2.
Forward-Time	Forward time is how long (in seconds) a port will spend in the listening-and-learning state before transitioning to forwarding state. Range of values is 4 to 30. Default value is 15.
Max-Age	The maximum age before the switch considers the received BPDU information on a port to be expired. Max-age is used when interworking with switches outside the region. Range of values is 6 to 40. Default value is 20.
Max-Hops	Maximum hops is used inside the MST region. Hop count is decremented each time the BPDU is forwarded. If max-hops reaches zero, the switch discards the BPDU and ages out the information on the receiving port. Range of values is 1 to 40. Default value is 20.

Using the CLI:

```
config switch stp settings
   set forward-time <4 - 30>
   set hello-time <1 - 10>
   set max-age <6 - 40>
   set max-hops <1 - 40>
   set name <region name>
   set revision <0 - x>
   set status {enable | disable}
end
```

Configuring an MST instance

STP topology is unique for each MST instance in the region. You can configure a different bridge priority and port parameters for each instance.

Using the Web-based manager:

1. Go to **Switch** > **STP** > **Instance.**
2. Create a new MST instance or select an existing instance to edit.
3. Update the instance parameters as described in the following table.
4. Click **Apply** to save the settings.

Settings	Guidelines
ID	Instance ID. Range is 1 - 15.
Priority	Priority is a component of bridge ID. The switch with the lowest bridge ID becomes the root switch for this MST instance. Allowed values: 0, 4096, 8192, 12288, 16384, 20480, 24576, 28672, 32768, 36864, 40960, 45056, 49152, 53248, 57344, and 61440. The default value is 32768.
VLAN Range	The VLANs that map to this MST instance. You can specify individual VLAN numbers or a range of numbers. **NOTE:** Do not assign any VLAN to more than one MST instance. Each VLAN number is in the range 1-4094.
Port Configuration	
Name	Port that will participate in this MST instance.
Cost	The switch uses port cost to select designated ports. Port cost is added to the received BPDU root cost in any BPDU sent on this port. A lower value is preferred. The range of values is 1 to 200,000,000. The default value depends on the interface speed: - 10 Gigabit Ethernet: 2,000 - Gigabit Ethernet: 20,000 - Fast Ethernet: 200,000 - Ethernet: 2,000,000
Priority	The switch uses port priority to choose among ports of the same cost. The port with the lowest priority is put into forwarding state. The valid values are: 0, 32, 64, 96, 128, 160, 192, and 224. The default value is 128.

Using the CLI:

```
config switch stp instance
```

```
edit <instance number>
   set priority <>
   config stp-port
      edit <port name>
         set cost <>
         set priority <>
      next
   set vlan-range <vlan range>
end
```

Example:

```
config switch stp instance
   edit "1"
      set priority 8192
      config stp-port
         edit "port18"
            set cost 0
            set priority 128
         next
         edit "port19"
            set cost 0
            set priority 128
         next
      end
      set vlan-range 5 7 11-20
   end
```

Configuring STP port settings

By default, STP (and edge port) is enabled on all ports.

Configuring an STP edge port

Use the following commands to enable or disable an interface as an STP edge port:

```
config switch interface
   edit port<number>
      set edge-port <enabled | disabled>
   next
end
```

Configuring STP loop protection

By default, STP loop protection is disabled on all ports. Use the following commands to configure STP loop protection on a port:

```
config switch interface
   edit port<number>
      set stp-loop-protection <enabled | disabled>
   next
end
```

Configuring STP root guard

Enable root guard on all ports that should not be root bridges. Do not enable root guard on the root port. You must have STP enabled to be able to use root guard.

To configure root guard on a port, use the following commands:

```
config switch interface
   edit port<number>
      set stp-root-guard <enable | disable>
   next
end
```

For example, to enable root guard on port 20:

```
config switch interface
   edit port20
      set stp-state enabled
      set stp-root-guard enable
   next
end
```

Configuring STP BPDU guard

There are three prerequisites for using BPDU guard:

- You must define the port as an edge port with the `set edge-port enabled` command.
- You must must enable STP on the switch interface with the `set stp-state enabled` command.
- You must enable STP on the global level with the `set status enable` command.

You can set how long the port will go down when a BPDU is received for a maximum of 120 minutes. The default port timeout is 5 minutes. If you set the timeout value to 0, the port will not go down when a BPDU is received, but you will have manually reset the port.

To configure BPDU guard on an STP edge port, use the following commands:

```
config switch interface
   edit port<number>
      set stp-bpdu-guard <enabled | disabled>
      set stp-bpdu-guard-timeout <0-120>
   next
end
```

For example, to enable BPDU guard on port 30 with a timeout value of 1 hour:

```
config switch stp settings
   set status enable
end
config switch interface
   edit port30
      set stp-state enabled
      set edge-port enabled
      set stp-bpdu-guard enabled
      set stp-bpdu-guard-timeout 60
   next
end
```

If you set the port timeout to 0, you will need to reset the port after it receives BPDUs and goes down. Use the following command to reset the port:

```
execute bpdu-guard reset port<number>
```

To check if BPDU guard has been triggered and on which ports, use the following command:

```
diagnose bpdu-guard display status
```

```
Portname             State      Status       Timeout(m)     Count    Last-Event
___________________  _________  ___________  _____________  ______  ____________________
port1                disabled     -              -            -          -
port2                disabled     -              -            -          -
port3                disabled     -              -            -          -
port4                disabled     -              -            -          -
port5                disabled     -              -            -          -
port6                disabled     -              -            -          -
port7                disabled     -              -            -          -
port8                disabled     -              -            -          -
port9                disabled     -              -            -          -
port10               disabled     -              -            -          -
port11               disabled     -              -            -          -
port12               disabled     -              -            -          -
port13               disabled     -              -            -          -
port14               disabled     -              -            -          -
port15               disabled     -              -            -          -
port16               disabled     -              -            -          -
port17               disabled     -              -            -          -
port18               disabled     -              -            -          -
port19               disabled     -              -            -          -
port20               disabled     -              -            -          -
port21               disabled     -              -            -          -
port22               disabled     -              -            -          -
port23               disabled     -              -            -          -
port25               disabled     -              -            -          -
port26               disabled     -              -            -          -
port27               disabled     -              -            -          -
port28               disabled     -              -            -          -
port29               disabled     -              -            -          -
port30               enabled      -              60           0          -
__FoRtI1LiNk0__      disabled     -              -            -          -
```

Interactions outside of the MSTP region

A boundary port on an MST switch is a port that receives an STP (version 0) BPDU, an RSTP (version 2) BPDU, or a BPDU from a different MST region.

If the port receives a version 0 BPDU, it will only send version 0 BPDUs on that port. Otherwise, it will send version 3 (MST) BPDUs because the RSTP switch will read this as an RSTP BPDU.

Viewing the MSTP configuration

To view the MSTP configuration details, use the following commands:

```
get switch stp instance
get switch stp settings
```

Use the following commands to display information about the MSTP instances in the network:

```
diagnose stp instance list
diagnose stp vlan list
diagnose stp mst-config list
```

Link aggregation groups

This chapter provides information on how to configure a link aggregation group (LAG). For LAG control, FortiSwitch supports the industry-standard Link Aggregation Control Protocol (LACP). FortiSwitch supports LACP in active and passive modes. In active mode, you can optionally specify the minimum and maximum number of active members in a trunk group.

FortiSwitch supports flap-guard protection for switch ports in a LAG.

This chapter covers the following topics:

Configuring the trunk and LAG ports

It is important to configure the trunk to prevent loops.

Using the Web-based manager:

1. Go to **Switch > Port > Trunk** and select **Create Trunk**.
2. Give the trunk an appropriate name.
3. Set **Mode** to **static**, **lacp-active**, or **lacp-passive.**
4. Add the required ports to the **Members** list.
5. Select **OK.**

Using the CLI:

```
config switch trunk
   edit <trunk name>
      set description <description_string>
      set members <ports>
      set mode {lacp-active | lacp-passive | static}
      set member-withdrawal-behavior {block | forward}
      set lacp-speed {fast | slow}
      set bundle [enable|disable]
         set min_bundle <integer>
         set max_bundle <integer>
      set port-selection-criteria
          {src-ip | src-mac | dst-ip |dst-mac | src-dst-ip |src-dst-mac}
   end
end
```

Example configuration

The following is an example CLI configurations for trunk/LAG ports:

Trunk/LAG ports

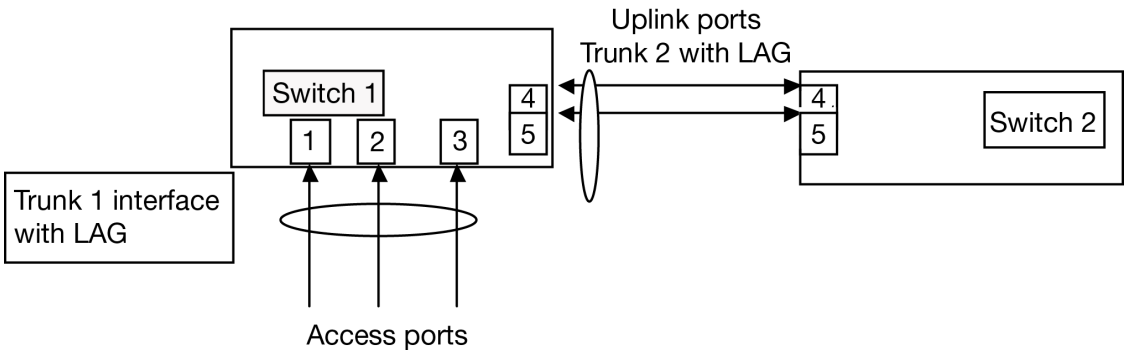

1. Configure the trunk 1 interface and assign member ports as a LAG group:

```
config switch trunk
   edit trunk1
      set members "port1" "port2" "port3"
      set description test
      set mode lacp-passive
      set port-selection criteria src-dst-ip
   end
end
```

2. Configure the switch ports to have native vlan assignments and allow those vlans on the port that will be the uplink port:

```
config switch interface
   edit port 1
      set native-vlan 1
   next
   edit port 2
      set native-vlan 2
   next
   edit port 3
      set native-vlan 3
   next
   edit port 4
      set native-vlan 4
      set allowed vlans 1 2 3
   next
   edit port 5
      set native-vlan 5
      set allowed-vlans 1 2 3
   end
end
```

3. Configure the trunk 2 interface and assign member ports as a LAG group:

```
config switch trunk
   edit trunk2
```

```
        set members "port4" "port5"
        set description test
        set mode lacp-passive
        set port-selection criteria src-dst-ip
    end
end
```

Checking the trunk configuration

To see the details of a configured trunk, use the following command:

```
diagnose switch trunk list
```

MCLAG

A link aggregation group (LAG) formed by an provides link-level redundancy. A multichassis LAG (MCLAG) provides node-level redundancy by grouping two FortiSwitch models together so that they appear as a single switch on the network. If either switch fails, the MCLAG continues to function without any interruption, increasing network resiliency and eliminating the delays associated with the Spanning Tree Protocol (STP).

This chapter covers the following topics:

Notes

- Fortinet recommends that both peer switches be of the same hardware model and same software version. Mismatched configurations might work but are unsupported.
- There is a maximum of two FortiSwitch models per MCLAG.
- The routing feature is not available within a MCLAG.
- Starting in FortiSwitchOS 3.6.4, by default, the MCLAG can use the STP.
- To use static MAC addresses within a MCLAG, you need to configure MAC addresses on both switches that form the LAG.

Example configuration

The following is an example CLI configurations for a MCLAG:

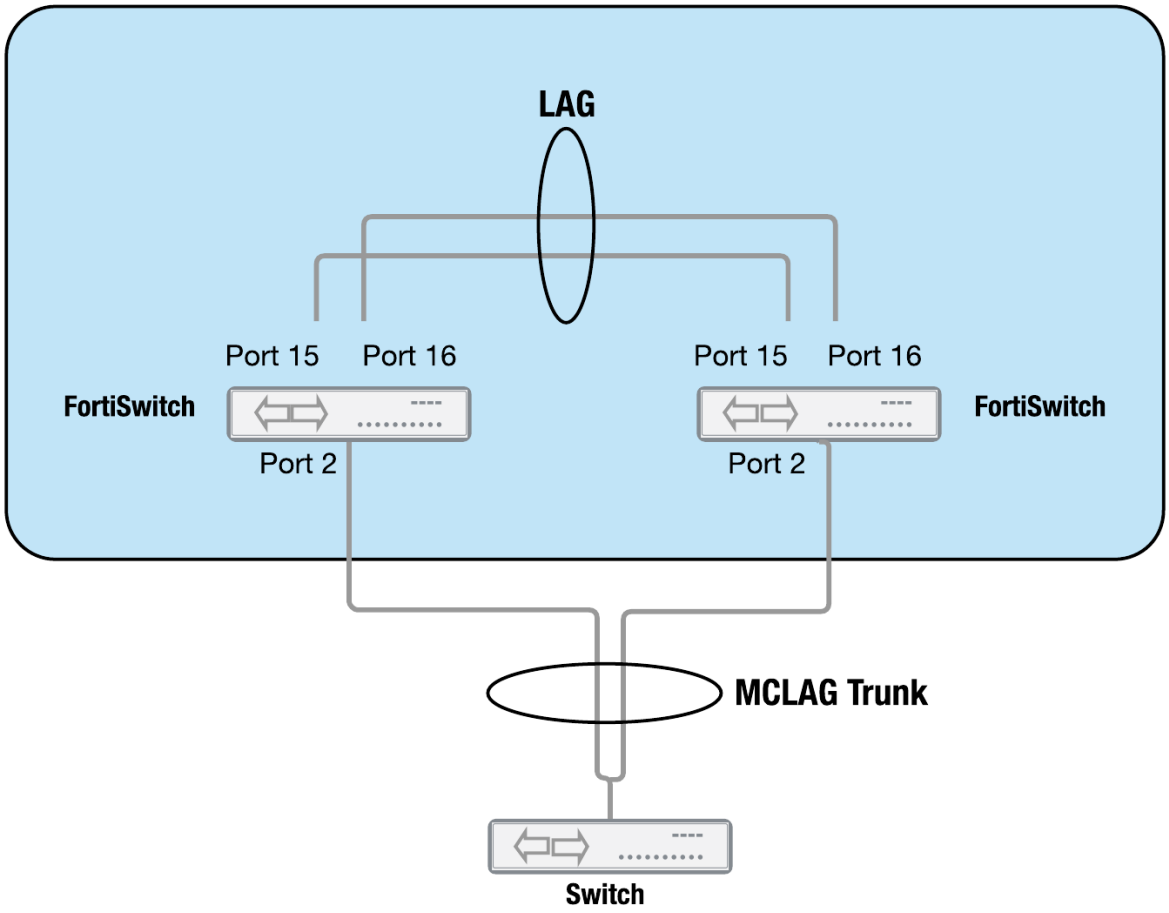

1. Create a LAG by configuring the ports for each FortiSwitch:

```
config switch trunk
   edit "MCLAG-ICL-trunk"
      set mclag-icl enable
      set members "port15" "port16"
      set mode lacp-active
   next
end
```

2. Set up the MCLAG:

```
config switch trunk
   edit "first-mclag"
      set mclag enable
      set members "port2"
   next
end
```

3. If you do not want the MCLAG to use the STP:

```
config switch global
   set mclag-stp-aware disabled
end
```

Viewing the configured trunk

To see the details of the MCLAG, use the following commands:

```
diagnose switch mclag icl
diagnose switch mclag list
```

Multi-stage load balance

You can use a FortiSwitch to configure multi-stage load balancing on a set of FortiGate units. This capability allows you to scale security processing while maintaining a simple basic architecture. This configuration is commonly referred to a “firewall sandwich.”

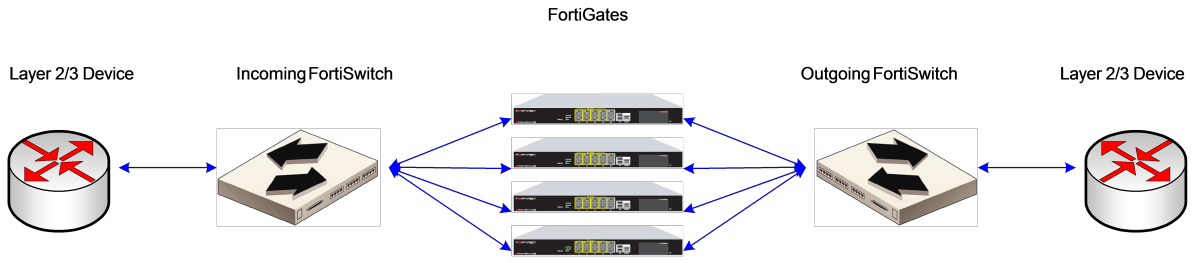

Because FortiGate provides session-aware analysis, the load distribution algorithm must be symmetric (traffic for a given session, in both directions, must all traverse the same FortiGate).

For larger scale deployment, the topology uses multiple layers of load distribution to allow for far larger numbers of FortiGate devices.

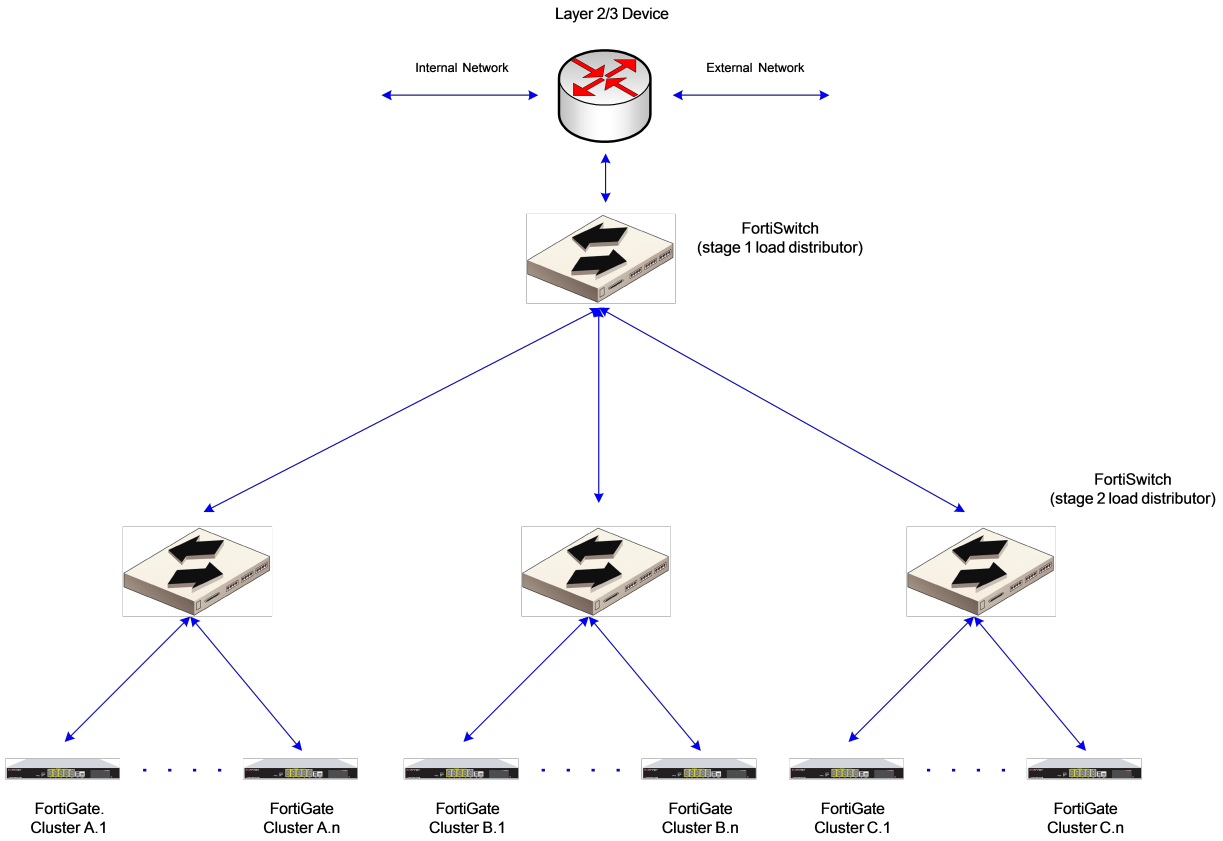

The hash at the first and second stages must be symmetric. The two stages must provide different hashing results.

This chapter covers the following topics:

Configuring the trunk ports

Use the following commands to configure the trunk members and set the port-selection criteria:

```
config switch trunk
   edit <trunk name>
      set description <description_string>
      set members <ports>
      set mode {fortinet-trunk | lacp-active | lacp-passive | static}
      set port-selection-criteria src-dst-ip-xor16
   end
end
```

Heartbeats

When in Fortinet-trunk mode, Heartbeat capability is enabled. Heartbeat messages monitor the status of FortiGate units. If one is unavailable, the FortiSwitch stops sending traffic to that FortiGate until the FortiGate becomes available.

If you enable `hb-verify`, each received heartbeat frame will be validated to match the signature (transmit-port plus switch serial number) and the following configured heartbeat parameters:

- hb-in-vlan
- hb-src-ip
- hb-dst-ip
- hb-src-upd-port
- hb-dst-udp-port

The destination MAC address of the heartbeat frame is set by default to 02:80:c2:00:00:02. You can change the value to any MAC address that is not a broadcast or multicast MAC address.

Configuring heartbeats

Configure the heartbeat fields using trunk configuration commands, as shown in this section. By default, all of the configurable values are set to zero, and hb-verify is disabled.

Set the mode to `forti-hb` and set the heartbeat loss limit to a value between 3 and 32.

The heartbeat will transmit at 1-second intervals on any link in the trunk that is up. This value is not configurable.

The heartbeat frame has configurable parameters for the layer-3 source and destination addresses and the layer-4 UDP ports. You must also specify the transmit and receive VLANs.

```
config switch trunk
   edit hb-trunk
      set mode fortinet-trunk
      set members <port> [<port>] ... [<port>]
      set hb-loss-limit <3-32>
      set hb-out-vlan <int>
      set hb-in-vlan <int>
      set hb-src-ip <x.x.x.x>
      set hb-dst-ip <x.x.x.x>
      set hb-src-udp-port <int>
      set hb-dst-udp-port <int>
      set hb-verify [ enable | disable ]
end
```

Use the following command to configure the destination MAC address:

```
config switch global
   set forti-trunk-dmac <mac address>
end
```

Example

The following example creates trunk tr1 with heartbeat capability:

```
config switch trunk
   edit "tr1"
      set mode fortinet-trunk
      set members "port1" "port2"
      set hb-out-vlan 300
      set hb-in-vlan 500
      set hb-src-ip 10.105.7.200
      set hb-dst-ip 10.105.7.199
      set hb-src-udp-port 12345
      set hb-dst-udp-port 54321
      set hb-verify enable
   next
end
```

LLDP-MED

The Fortinet data center switches support the Link Layer Discovery Protocol (LLDP) for transmission and reception wherein the switch will multicast LLDP packets to advertise its identity and capabilities. A switch receives the equivalent information from adjacent layer-2 peers.

Fortinet data center switches support LLDP-MED (Media Endpoint Discovery), which is an enhancement of LLDP that provides the following facilities:

- Auto-discovery of LAN policies (such as VLAN, layer-2 priority, and differentiated services settings), to enable plug-and-play networking.
- Device location discovery to allow the creation of location databases and Enhanced 911 services for Voice over Internet Protocol (VoIP).
- Extended and automated power management for power over Ethernet (PoE) endpoints.
- Inventory management, allowing network administrators to track their network devices, and determine their characteristics (manufacturer, software and hardware versions, serial or asset number).

The switch will multicast LLDP packets to advertise its identity and capabilities. The switch receives the equivalent information from adjacent layer-2 peers.

This chapter covers the following topics:

Configuration notes

Fortinet recommends LLDP-MED-capable phones.

The FortiSwitch functions as a Network Connectivity device (that is, NIC, switch, router, and gateway), and will only support sending TLVs intended for Network Connectivity devices.

LLDP supports up to 16 neighbors per physical port.

FortiSwitch accepts and parses packets using the CDP (Cisco Discovery Protocol) and count CDP neighbors towards the neighbor limit on a physical port. If neighbors exist, FortiSwitch transmits CDP packets in addition to LLDP.

With release 3.5.1, CDP is independently controllable through **cdp-status** on the physical port. The FortiSwitch no longer requires a neighbor to trigger it to transmit CDP; it will transmit provided cdp-status is configured as tx-only or tx-rx. The default configuration for CDP-status is disabled. It still uses values pulled from the lldp-profile to configure its contents.

LLDP must be globally enabled in switch.lldp.settings for CDP to be transmitted or received:

NOTE: If a port is added into a *virtual-wire* (connects two ends of a controlled system using a radio frequency [RF] medium), the FortiSwitch will disable transmit and receipt of LLDP and CDP packets and remove all neighbors from the port. This virtual-wire state is noted in the `get switch lldp neighbor-summary` command output.

If the combination of configured TLVs exceeds the maximum frame size on a port, that frame cannot be sent.

LLDP global settings

Using the Web-based manager:

1. Go to **Switch > LLDP MED > Settings**.
2. Enable or disable the status.
3. Enter a value for TX hold.
4. Enter the number of seconds for TX interval.
5. Enable or disable the fast start; if you enable fast start, enter the number of seconds.
6. Select the management interface.
7. Click **OK**.

Using the CLI:

```
config switch lldp settings
   set status < enable | disable >
   set tx-hold <int>
   set tx-interval <int>
   set fast-start-interval <int>
   set management-interface <layer-3 interface>
end
```

Variable	Description
status	Enable or disable
tx-hold	Number of tx-intervals before the local LLDP data expires (that is, the packet TTL (in seconds) is **tx-hold** times **tx-interval**). The range for tx-hold is 1 to 16, and the default value is 4.
tx-interval	Frequency of LLDP PDU transmission ranging from 5 to 4095 seconds (default is 30).
fast-start-interval	How often the FortiSwitch transmits the first four LLDP packets when a link comes up. The range is 2 to 5 seconds and the default is 2 seconds. Set this variable to zero to disable fast start.
management-interface	Primary management interface advertised in LLDP and CDP PDUs.

Setting the asset tag

To help identify the unit, LLDP uses the asset tag, which can be at most 32 characters. It will be added to the LLDP-MED inventory TLV (when that TLV is enabled):

```
config system global
   set asset-tag <string>
end
```

Configuring LLDP profiles

LLDP profile contains most of the port-specific configuration. Profiles are designed to provide a central point of configuration for LLDP settings that are likely to be the same for multiple ports.

Two static LLDP profiles, default and default-auto-isl, are created automatically. They can be modified but not deleted. The default-auto-isl profile always has auto-isl enabled and rejects any configurations that attempt to disable it.

LLDP-MED network policies

LLDP-MED network policies cannot be deleted or added. To use a policy, set the med-tlvs field to include `network-policy` and the desired network policy to `enabled`. The VLAN values on the policy are cross-checked against the VLAN native and untagged attributes for any interfaces that contain physical-ports using this profile. The cross-check determines if the policy Type Length Value (TLV) should be sent (VLAN must be native or allowed) and if the TLV should mark the VLAN as tagged or untagged (VLAN is native, or is in untagged). The network policy TLV is automatically updated when either a switch interface changes VLAN configuration or a physical port is added to, or removed from, a trunk.

FortiSwitch will support the following LLDP-MED TLVs:

- Network Policy TLV
- Inventory Management TLV

Refer to the Configuration deployment example on page 94.

Custom TLVs (organizationally specific TLVs)

Custom TLVs are configured in their own subtable, available in each profile. They allow you to emulate the TLVs defined in various specifications by using their OUI and subtype and ensuring that the data is formatted correctly. You could also define a purely arbitrary custom TLV for some other vendor or for their company.

The "name" value for each custom TLV is neither used by nor has an effect on LLDP; it simply differentiates config entries:

```
config custom-tlvs
edit <name>
```

The OUI value for each TLV must be set to three bytes. If just one of those bytes is nonzero it is accepted; any value other than "000" is valid. The subtype is optional and ranges from 0 (default) to 255. The information string can be 0 to 507 bytes, in hexadecimal notation.

FortiSwitch does not check for conflicts either between custom TLV values or with standardized TLVs. That is, other than ensuring that the OUI is nonzero, FortiSwitch do not check the OUI, subtype (or data) values entered in the CLI for conflicts with other Custom TLVs or with the OUI and subtypes of TLVs defined by the 802.1, 802.3, LLDP-MED, or other standards. While this behavior could cause LLDP protocol issues, it also allows a large degree of flexibility were you to substitute a standard TLV that is not supported yet.

802.1 TLVs

The only 802.1 TLV that can be enabled or disabled is port-vlan-id. This TLV will send the native VLAN of the port. This value is updated when the native VLAN of the interface representing the physical port changes or if the physical port is added to, or removed from, a trunk.

By default, no 802.1 TLVs are enabled.

802.3 TLVs

The only 802.3 TLV that can be enabled or disabled is max-frame-size. This TLV will send the max-frame-size value of the port. If this variable is changed, the sent value will reflect the updated value.

By default, no 802.3 TLVs are enabled.

Auto-ISL

The auto-ISL configuration that was formerly in the `switch physical-port` command has been moved to the `switch lldp-profile` command. All behavior and default values are unchanged.

Configuring an LLDP profile for the port

Configure an LLDP profile for the port. By default, the port uses the default LLDP profile.

Using the Web-based manager:

1. Go to **Switch > LLDP MED > Profile**.
2. Click **Create New**.
3. Enter a name for your LLDP profile.
4. If needed, select port-vlan-id.
5. If needed, select max-frame-size.
6. If needed, enable auto-isl.
7. Enter a value for the auto-isl-hello-timer.
8. Enter a value for the auto-isl-port-group.
9. Enter a value for the auto-isl-receive-timeout.
10. If needed, select inventory-management, network-policy, or both.
11. Click **OK**.

Using the CLI:

```
config switch lldp profile
   edit <profile>
```

```
set 802.1-tlvs port-vlan-id
set 802.3-tlvs max-frame-size
set auto-isl {active | inactive}
set auto-isl-hello-timer <1-30>
set auto-isl-port-group <0-9>
set auto-isl-receive-timeout <3-90>
set med-tlvs (inventory-management | network-policy)
```

Enabling LLDP on a port

To enable LLDP MED on a port, set the LLDP status to receive-only, transmit-only, or receive and transmit. The default value is tx-rx.

Using the Web-based manager:

1. Go to **Switch > Port > Physical**.
2. Select a port and click **Edit**.
3. Select **tx rx**, **rx only**, **tx only**, or **disable**.
4. Select an LLDP profile.
5. Click **OK**.

Using the CLI:

```
config switch physical-port
   edit <port>
      set lldp-status (rx-only | tx-only | tx-rx | disable)
      set lldp-profile <profile name>
   next
end
```

Checking the LLDP configuration

View the LLDP configuration settings using the Web-based manager:

1. Go to **Switch > LLDP MED > Settings**.
2. Click **OK**.

View the LLDP configuration settings using the CLI:

```
get switch lldp settings
status : enable
tx-hold : 4
tx-interval : 30
fast-start-interval : 2
management-interface: internal
```

View the LLDP profiles using the Web-based manager:

1. Go to **Switch > LLDP MED > Profile**.
2. Select a profile and click **Edit**.
3. Click **OK**.

View the LLDP profiles using the CLI:

```
get switch lldp profile
== [ default ]
name: default 802.1-tlvs: 802.3-tlvs: med-tlvs: inventory-management network-policy
== [ default-auto-isl ]
name: default-auto-isl 802.1-tlvs: 802.3-tlvs: med-tlvs:
```

Use the following commands to display the LLDP information about LLDP status or the layer-2 peers for this FortiSwitch:

```
get switch lldp (auto-isl-status | neighbors-detail | neighbors-summary | profile |
    settings | stats)
```

Configuration deployment example

Configuring LLDP includes the following steps:

1. Configure LLDP global configuration settings using the `config switch lldp settings` command.
2. Create LLDP profiles using the `config switch lldp profile` command to configure Type Length Values (TLVs) and other per-port settings. (TLVs)
3. Assign LLDP profiles to physical ports.
4. Apply VLAN to interface. (**NOTE:** LLDP profile values that are tied to VLANs will only be sent if the VLAN is assigned on the switch interface.)
   a. Configure profile.

```
show switch lldp profile Forti670i
config switch lldp profile
   edit "Forti670i"
      config med-network-policy
         edit "voice"
            set dscp 46
            set priority 5
            set status enable
            set vlan 400
         next
         edit "guest-voice"
         next
         edit "guest-voice-signaling"
         next
         edit "softphone-voice"
         next
         edit "video-conferencing"
         next
```

```
            edit "streaming-video"
               set dscp 40
               set priority 3
               set status enable
               set vlan 400
            next
            edit "video-signalling"
            next
         end
      set med-tlvs inventory-management network-policy
   next
   end
```

b. Configure the interface.

```
show switch interface port4
config switch interface
   edit "port4"
      set allowed-vlans 400
      set snmp auto
   next
end
```

c. Connect a phone with LLDP-MED capability to the interface. **NOTE:** Make certain the LLDP, Learning, and DHCP features are enabled.

```
show switch physical-port port4
config switch physical-port
   edit "port4"
      set lldp-profile "Forti670i"
      set speed auto
   next
end
```

d. Verify.

```
show switch lldp neighbor-det port4

Neighbor learned on port port4 by LLDP protocol
Last change 12 seconds ago
Last packet received 12 seconds ago
Chassis ID: 10.105.251.40 (ip)
System Name: FON-670i
System Description:
V12.740.335.12.B
Time To Live: 60 seconds
System Capabilities: BT
Enabled Capabilities: BT
MED type: Communication Device Endpoint (Class III)
MED Capabilities: CP
Management IP Address: 10.105.251.40
Port ID: 00:a8:59:d8:f1:f6 (mac)
Port description: WAN Port 10M/100M/1000M
IEEE802.3, Power via MDI:
Power devicetype: PD
PSE MDI Power: Not Supported
```

```
PSE MDI Power Enabled: No
PSE Pair Selection: Can not be controlled
PSE power pairs: Signal
Power class: 1
Power type: 802.3at off
Power source: Unknown
Power priority: Unknown
Power requested: 0
Power allocated: 0
LLDP-MED, Network Policies:
voice: VLAN: 400 (tagged), Priority: 5 DSCP: 46
voice-signaling: VLAN: 400 (tagged), Priority: 4 DSCP: 35
streaming-video: VLAN: 400 (tagged), Priority: 3 DSCP: 40
```

Checking LLDP details

Using the Web-based manager:

1. Go to **Switch > Monitor > LLDP**.
2. Select an entry.
3. Click **display detail**.

MAC/IP/protocol-based VLANs

The FortiSwitch assigns VLANs to packets based on the incoming port or the VLAN tag in the packet. The MAC/IP/protocol-based VLAN feature enables the assignment of VLANs based on specific fields in an ingress packet (MAC address, IP address, or layer-2 protocol).

This chapter covers the following topics:

Overview

When a MAC/IP/protocol-based VLAN is assigned to a port, the default behavior is for egress packets with that VLAN value to include the VLAN tag. Use the `set untagged-vlans <vlan>` configuration command to remove the VLAN tag from egress packets. For an example of the command, see the Example configuration on page 99.

The MAC/IP/protocol-based VLAN feature assigns the VLAN based on MAC address, IP address, or layer-2 protocol.

MAC based

In MAC-based VLAN assignment, the FortiSwitch associates a VLAN with each packet based on the originating MAC address.

IP based

In IP-based VLAN assignment, the FortiSwitch associates a VLAN with each packet based on the originating IP address or IP subnet. IPv4 is supported with prefix masks from 1 to 32. IPv6 is also supported, depending on hardware availability, with prefix lengths from 1 to 64.

Protocol based

In protocol-based VLAN assignment, the FortiSwitch associates a VLAN with each packet based on the Ethernet protocol value and the frame type (ethernet2, 802.3d/SNAP, LLC).

Configuring MAC/IP/protocol-based VLANs

Note the following prerequisites:

- The VLAN must be created in the FortiSwitch
- The VLAN needs to be allowed on the ingress port

Using the Web-based manager:

1. Go to **Switch** > **VLAN** > **MAC/IP Membership**.
2. Click **Create New** for a new VLAN or select a VLAN and click **Edit**.
3. To configure a MAC-based VLAN:
   a. Click **New** under Member By MAC.
   b. Enter a description and the MAC address.
   c. To save the entry, click the plus icon (**+**) to the right of the new entry.
4. To configure an IP-based VLAN:
   a. Click **New** under Member By IPV4.
   b. Enter a description and the IP and Mask.
   c. To save the entry, click the plus icon (**+**) to the right of the new entry.
5. Click **OK**.

Using the CLI:

```
config switch vlan
   edit <vlan-id>
      config member-by-mac
         edit <id>
            set mac xx:xx:xx:xx:xx:xx
            set description <128 byte string>
         next
      end
      config member-by-ipv4
         edit <id>
            set address a.b.c.d/e #subnet mask must 1-32
            set description <128 byte string>
         next
      end
      config member-by-ipv6
         edit <id>
            set prefix xx:xx:xx:xx::/prefix #prefix must 1-64
            set description <128 byte string>
         next
      end
      config member-by-proto
         edit <id>
            set frametypes ethernet2 802.3d llc #default is all
            set protocol 0xXXXX
         next
      end
   next
end
```

NOTE: There are hardware limits regarding how many MAC/IP/Protocol-based VLANs.you can configure. If you try to add entries beyond the limit, the CLI will reject the configuration:

- editing an existing VLAN - when you enter **next** or **end** on the **config member-by** command
- adding a new VLAN - when you enter **next** or **end** on the **edit vlan** command

Example configuration

The following example shows a CLI configuration for MAC-based VLAN where a VOIP phone and a PC share the same switch port.

In this example, a unique VLAN is assigned to the voice traffic, and the PC traffic is on the default VLAN for the port.

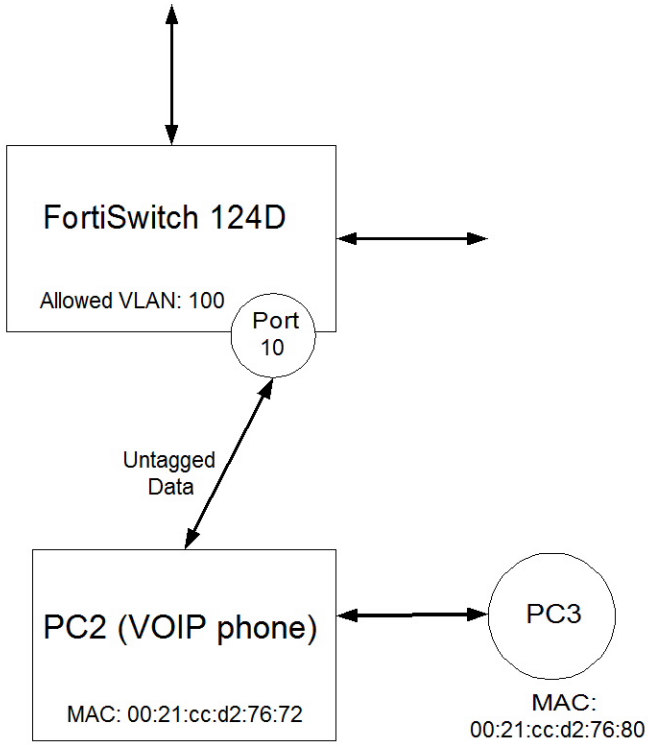

1. Switch FS-124D Port 10 is connected to PC2 (a VOIP phone), with MAC address 00:21:cc:d2:76:72.
2. The phone also sends traffic from PC3 (MAC= 00:21:cc:d2:76:80).
3. Assign the PC3 traffic to the default VLAN (1) on port 10.
4. Assign the voice traffic to VLAN 100.

Configure the voice VLAN

```
config switch vlan
   edit 100
      config member-by-mac
         edit 1
            set description "pc2"
            set mac 00:21:cc:d2:76:72
         next
      end
   end
end
```

Configure switch port 10

```
config switch interface
   edit "port10"
      # allow vlan=100 on this port
      # treat this as untagged on egress
      set allowed-vlans 100
      set untagged-vlans 100
      set snmp-index 10
   end
end
```

Checking the configuration

To view the MAC-based VLAN assignments, use the following command:

```
diagnose switch vlan assignment mac list sorted-by-mac

  00:21:cc:d2:76:72   VLAN: 100 Installed: yes
 Source: Configuration (entry 1)
 Description: pc2
```

Mirroring

This chapter contains information on how to configure layer-2 port mirroring.

The following topics are covered:

Configuring a mirror

NOTE: You can use virtual wire ports as ingress and egress mirror sources. Egress mirroring of virtual wire ports will have an additional VLAN header on all mirrored traffic.

Using the Web-based manager:

1. Go to **Switch** > **Mirror** > **Mirror**.
2. Click **Create New.**
3. Enter a name for the mirror.
4. Set the **Status Enable** check box to set the mirror to active.
5. Select a Destination Port.
6. Select available ports to be used for Ingress Monitoring and Egress Monitoring.
7. Enable the **Packet switching functionality when mirroring** option if the destination port is not a dedicated port. For example, enable this option if you connect a laptop to the switch and you are running a packet sniffer along with the management GUI on the laptop:

Using the CLI:

```
config switch mirror
   edit "m1"
      set dst "port5"
      set src-egress "port2" "port3"
      set src-ingress "port2" "port4"
      set status active
      set switching-packet enable
end
```

Multiple mirror destination ports (MTPs)

With some FortiSwitch models, you can configure multiple mirror destination ports with the following guidelines and restrictions:

- Always set the destination port before setting the src-ingress or src-egress ports.
- Any port configured as a src-ingress or src-egress port in one mirror cannot be configured as a destination port in another mirror.

- For switch models FS-1024D, FS-1048D, FS-3032D, and FS-5xx Series:
  - You can configure a maximum of four mirror destination ports.
  - Multiple ingress or egress ports can be mirrored to the same destination port.
  - The same ingress/egress port can be mirrored to more than one destination port.
- For switch models FSW-124D, FS-124D-POE, FS-224D-FPOE, FS-2xxE Series, and FS-4xxD Series:
  - You can configure a maximum of four mirror destination ports.
  - Multiple ingress or egress ports can be mirrored to the same destination port.
  - A source ingress port cannot be mirrored to more than one destination port.
  - All source egress ports must be mirrored to the same destination port.
- For switch models FS-108D-POE and FSR-112D-POE:
  - You can configure up to seven mirrors, each with a different destination port.
  - Multiple ingress or egress ports can be mirrored to the same destination port.
  - An ingress or egress port cannot be mirrored to more than one destination port.

These restrictions apply to active mirrors. If you try to activate an invalid mirror configuration, the system will display the `Insufficient resources!!` error message.

The following example configuration is valid for FortiSwitch-3032D. This configuration includes three ingress ports, one egress port, and four destination ports. The port3 ingress and egress ports are mirrored to multiple destinations.

```
config switch mirror
   edit "m1"
      set dst "port16"
      set status active
         set src-ingress "port3" "port5" "port7"
   next
   edit "m2"
      set dst "port22"
      set status active
         set src-ingress "port3" "port5"
   next
   edit "m3"
      set dst "port1"
      set status active
         set src-egress "port3"
   next
   edit "m4"
      set dst "port2"
      set status active
         set src-egress "port3"
end
```

The following example configuration includes three ingress ports, three egress ports and four destination ports. Each ingress and egress port is mirrored to only one destination port.

```
config switch mirror
   edit "m1"
      set dst "port1"
      set status active
         set src-ingress "port2" "port7"
   next
   edit "m2"
      set dst "port5"
```

```
            set status active
               set src-egress "port2"
         next
         edit "m3"
            set dst "port3"
            set status active
               set src-ingress "port6"
         next
         edit "m4"
            set dst "port4"
            set status active
               set src-egress "port6" "port8"
      end
```

Access control lists

You can use access control lists (ACLs) to configure policies for different types of incoming traffic.

This chapter covers the following topics:

ACL policy attributes

Key attributes of a policy include:

- **Interface**. The interface(s) on which traffic arrives at the switch. The interface can be a port, a trunk, or all interfaces. The policy applies to ingress traffic only (not egress traffic).
- **Classifier**. The classifier identifies the packets that the policy will act on. Each packet can be classified based on one or more criteria. Criteria include source and destination MAC address, VLAN id, source and destination IP address, or service (layer 4 protocol id and port number).
- **Marking** involves setting bits in the packet header to indicate the priority of this packet.
- **Actions**. If a packet matches the classifier criteria for a given ACL, the following types of action may be applied to the packet:
    - allow or block the packet, redirect the packet, mirror the packet
    - police the traffic
    - mirror the packet to another port, interface or trunk
    - CoS queue assignment
    - outer VLAN tag assignment
    - egress mask to filter packets

The switch uses specialized TCAM memory to perform ACL matching. Each model of FortiSwitch provides different ACL-related capabilities. When you configure the ACL policy, the system will reject the request if the hardware cannot support it.

Configuring an ACL policy

Major steps to configure an ACL policy include the following:

1. (Optional) Create or customize a service. FortiSwitch provides a set of pre-configured services that you can use. Use the following command to list the services:

```
show switch acl service custom
```

2. (Optional) Create a policer, if you are defining ACLs to police different types of traffic.
3. Configure the security policies.

Details for each step are as follows:

1. (Optional) Create or customize a service:

```
config switch acl service custom
   edit <service name>
      set comment <string>
      set color
      set protocol {ICMP | IP | TCP/UDP/SCTP}
      set sctp-portrange <dstportlow_int>[-<dstporthigh_int>: <srcportlow_int>-
          <srcporthigh_int>]
      set tcp-portrange <dstportlow_int>[-<dstporthigh_int>:
      <srcportlow_int>-<srcporthigh_int>]
      set udp-portrange <dstportlow_int>[-<dstporthigh_int>:<srcportlow_int>-
          <srcporthigh_int>]
   end
end
```

2. (Optional) Create a policer:

```
config switch acl policer
   edit <policer index>
      set description
      set guaranteed-bandwidth <bandwidth_value>
      set guaranteed-burst <in_bytes>
      set maximum-burst <in_bytes>
```

   Each policy is assigned a unique policy ID that is automatically assigned. To view it, use the `get switch acl policy` command.

3. Configure the policy:

```
config switch acl policy
edit <policy-id>
   set description
   set ingress-interface < port >
   set ingress-interface-all {enable | disable}
   config classifier
      set src-mac <mac> <mask>
      set dst-mac <mac> <mask>
      set ether-type <integer>
      set src-ip-prefix <IP address> <mask>
      set dest-ip-prefix <IP address> <mask>
      set service <service-id>
      set vlanid <vlan-id>
      set cos <802.1Q CoS value to match>
      set dscp <DSCP value to match>
   end
   config action
      set count {enable | disable}
      set drop {enable | disable}
      set mirror [internal | <port> | <interface> | <trunk>]
      set outer-vlan-tag <integer>
      set policer <policer>
      set redirect [internal | <port>]
      set redirect-bcast-cpu {enable | disable}
      set redirect-bcast-no-cpu {enable | disable}
```

```
            set redirect-physical-port <port>
            set remark-cos <0-7>
            set remark-dscp <0-63>
         end
      end
```

Egress mask

Use the following commands to prevent specific ports from being used for egress:

```
config switch acl policy
   edit <policy-id>
      config classifier
   end
   config action
      set egress-mask <list of physical ports>
   end
end
```

NOTE: The egress-mask command is not supported on dual-chip platforms, such as 448D, 448D-POE, and 448D-FPOE.

Viewing counters

Use the following command to display the counters associated with a policy. If you do not provide a policy identifier, the system displays all policies that have counters:

```
get switch acl counters [policy-id]
```

```
ID      Packets                 Bytes            description
_____________________________________________________________

0001 1861642              119145728             ip_mac_filter
0100 11160319             714260416             udp_vlan_filter
```

Clearing counters

Use the following command to clear the counters associated with a policy. If you do not provide a policy ID, the system clears all of the ACL counters:

```
execute acl clear-counter <policy-id>
```

Configuration examples

Example 1

In the following example, traffic from VLAN 3 is blocked to a specified destination IP subnet (10.10.0.0/16) but allowed to all other destinations:

```
config switch acl policy
   edit 1
      config action
```

```
        set count enable
        set drop enable
      end
      config classifier
        set dst-ip-prefix 10.10.0.0 255.255.0.0
        set vlan-id 3
      end
      set ingress-interface-all enable
    next
    edit 2
      config classifier
        set vlan-id 3
      end
      set ingress-interface-all enable
    next
end
```

Example 2

In the following example, Server Message Block (SMB) traffic received on port 1 is mirrored to port 3. SMB protocol uses port 445:

```
config switch acl service custom
  edit "SMB"
    set tcp-portrange 445
  next
end
config switch acl policy # apply policy to port 1 ingress and send to port 3
  edit 1
    set description "cnt_n_mirror_smb"
    set ingress-interface "port1"
    config action
      set count enable
      set mirror "port3"
    end
    config classifier
      set service "SMB"
      set src-ip-prefix 20.20.20.100 255.255.255.255
      set dst-ip-prefix 100.100.100.0 255.255.255.0
    end
  next
end
```

Example 3

FortiSwitch can map different flows (for example, based on source and destination IP addresses) to specific outgoing ports.

In the following example, flows are redirected (based on destination IP) to different outgoing ports, connected to separate FortiDDOS appliances. This allows you to apply different FortiDDOS service profiles to different types of traffic:

```
config switch acl policy # apply policy to port 1 ingress and send to port 3
  edit 1
    config action
      set count enable
```

```
      set redirect "port3‟ # use redirect to shift selected traffic to new destination
    end
    config classifier
      set dst-ip-prefix 100.100.100.0 255.255.255.0
    end
    set description "cnt_n_mirror13"
    set ingress-interface "port1"
  next
  edit 2
    config action # apply policy to port 3 ingress and send to port 1
      set count enable
      set redirect "port1"
    end
    config classifier
      set src-ip-prefix 100.100.100.0 255.255.255.0
    end
    set description "cnt_n_mirror31"
    set ingress-interface "port3"
  next
end

config switch acl policy # apply policy to port 1 ingress and send to port 4
  edit 3
    config action
      set count enable
      set redirect "port4‟ # use redirect to shift selected traffic to new destination
    end
    config classifier
      set dst-ip-prefix 20.20.20.0 255.255.255.0
    end
    set description "cnt_n_mirror14"
    set ingress-interface "port1"
  next
  edit 4
    config action # apply policy to port 4 ingress and send to port 1
      set count enable
      set redirect "port1"
    end
    config classifier
      set src-ip-prefix 20.20.20.0 255.255.255.0
    end
    set description "cnt_n_mirror41"
    set ingress-interface "port4"
  next
end
```

Storm control

Storm control protects a LAN from disruption by traffic storms, which stem from mistakes in network configuration or denial-of-service attacks. A traffic storm, which may consist of broadcast, multicast, or unicast traffic, creates excessive traffic on the LAN and degrades network performance.

By default, storm control is disabled on a FortiSwitch. When enabled, it measures the data rate (in packets-per-second) for unknown unicast, unknown multicast, and broadcast traffic.

You can enable and disable storm control for each of these traffic types individually. If the traffic rate for any of the types exceeds the configured threshold, FortiSwitch drops the excess traffic.

Storm control configuration is global.

This chapter covers the following topics:

Configuring storm control

If you set the rate to zero, the system drops all packets (for the enabled traffic types):

Using the Web-based manager:

1. Go to **Switch > Storm Control > Settings**.
2. Enable **Broadcast**, **Unknown Unicast**, and **Unknown Multicast** as required.
3. Enter a value for the rate.
4. Click **Apply** to save the changes.

Using the CLI:

Use the following commands to configure storm control:

```
config switch storm-control
   set rate [0 | 1 - 100000]
   set unknown-unicast {enable | disable}
   set unknown-mcast {enable | disable}
   set broadcast {enable | disable}
```

Displaying the storm-control configuration

Use the following command to display the storm-control configuration:

```
get switch storm-control
```

DHCP snooping

The DHCP snooping feature monitors the DHCP traffic from untrusted sources (for example, typically host ports and unknown DHCP servers) that might initiate traffic attacks or other hostile actions. To prevent this, DHCP snooping filters messages on untrusted ports by performing the following activities:

- Validating DHCP messages received from untrusted sources and filtering out invalid messages. For example, a request to decline an DHCP offer or release a lease is ignored if the request is from a different interface than the one that created the entry.
- Building and maintaining a DHCP snooping binding database, which contains information about untrusted hosts with leased IP addresses.

Other security features like dynamic ARP inspection (DAI), a security feature that rejects invalid and malicious ARP packets, also use information stored in the DHCP snooping binding database.

In the FortiSwitch, all ports are untrusted by default, and DHCP snooping is disabled on all untrusted ports. You indicate that a source is trusted by configuring the trust state of its connecting interface.

FortiSwitch supports the option of including option-82 data in the DHCP request. (DHCP option 82 provides additional security by enabling a controller to act as a DHCP relay agent to prevent DHCP client requests from untrusted sources.)

For the option-82 Circuit ID field, the following format is used:

Circuit-ID: vlan-mod-port

mod - [(1 Byte) -> Snoop - 1 , Relay - 0]

vlan - [2 bytes]

port - [1 byte]

For the option-82 Remote ID field, the following format is used:

Remote-ID: mac [6 byte]

This chapter covers the following topics:

Configuring DHCP snooping

DHCP snooping is enabled per VLAN and, by default, DHCP snooping is disabled.

Configuring DHCP snooping consists of the following steps:

1. Configure the VLAN settings.
2. Configure the interface settings.

Configure VLAN settings

Using the Web-based manager:

1. Go to **Switch > VLAN > DHCP Snooping**.
2. Click **Create New**.
3. Enter the VLAN identifier.
4. Enter a description for the new VLAN.
5. Enable or disable the private VLAN.
6. If you enabled the private VLAN, enter the number of isolated subVLANs and enter which subVLANS belong to the community, separated by commas.
7. Select the **DHCP Snooping** check box.
8. If needed, select the **Verify Source Mac**, **Insert Option-82**, and **Dynamic ARP Inspection** check boxes.
9. If needed, select the **IGMP Snooping** check box.
10. Click **OK**.

Using the CLI:

```
config switch vlan
   edit <vlan-id>
      set dhcp-snooping <enable | disable>
      set dhcp-snooping-verify-mac <enable | disable>
      set dhcp-snooping-option82 <enable | disable>
   next
end
```

For example:

```
config switch vlan
   edit 10
      set dhcp-snooping enable
      set dhcp-snooping-verify-mac enable
      set dhcp-snooping-option82 enable
   next
end
```

NOTE: If you enable `dhcp-snooping-verify-mac`, the system will verify that the source MAC address in the DHCP request from an untrusted port matches the client hardware address.

NOTE: If you enable `dhcp-snooping-option82`, the system inserts option-82 data into the DHCP messages for this VLAN.

Configure interface settings

After you enable DHCP snooping on a VLAN, all interfaces are in an untrusted state by default, and DHCP snooping is disabled on all untrusted interfaces. You must explicitly configure the trusted interfaces and enable DHCP snooping for each interface.

In addition, you can set a limit for how many IP addresses are in the DHCP snooping binding database for each interface by enabling the `dhcp-snoop-learning-limit-check` and setting the `learning-limit`. By

default, `dhcp-snoop-learning-limit-check` is disabled, and the number of entries for an untrusted ports is 5. You can set the number of entries to 0. The maximum number of entries depends on which FortiSwitch you are using. For example:

```
S548DN4K16000313 # show switch vlan 1
config switch vlan
   edit 1
      set learning-limit 100
      set dhcp-snooping enable
   next
end
```

NOTE: If the FortiSwitch has already learned more IP addresses than the `dhcp-snoop-learning-limit` before the limit is set, the configuration is rejected because the FortiSwitch cannot select which IP addresses should be kept. If the FortiSwitch has learned fewer IP address or the same number of IP addresses as the `dhcp-snoop-learning-limit` before the limit is set, the configuration is accepted.

NOTE: The per-VLAN learning limit is not supported on dual-chip platforms (248 and 448 series)

Using the Web-based manager:

1. Go to **Switch > Interface > Interface** or **Switch > Interface > Trunk**.
2. Select an interface.
3. Click **Edit**.
4. Select a **Trusted** or **Untrusted** interface for DHCP snooping.
5. If you want to accept DHCP messages with option-82 data from an untrusted interface, select the **Option-82 Trust** check box.
6. Click **OK**.

Using the CLI:

```
config switch interface / trunk
   edit <interface-name>
      set native-vlan <VLAN-ID>
      set dhcp-snooping {trusted | untrusted}
      set dhcp-snoop-learning-limit-check {enable | disable}
      set learning-limit <integer>
      set dhcp-snoop-option82-trust {enable | disable}
   next
end
```

For example:

```
config switch interface
   edit "port5"
      set native-vlan 10
      set dhcp-snooping untrusted
      set dhcp-snoop-learning-limit-check enable
      set learning-limit 7
      set dhcp-snoop-option82-trust enable
      set snmp-index 5
   next
end
```

Set `dhcp-snooping` to reflect the trust state of the interface. Where DHCP servers are located, you must configure interfaces as trusted.

If you enable `dhcp-snoop-option82-trust`, the system accepts DHCP messages with option-82 data from an untrusted interface.

Checking the DHCP snooping configuration

Use the following command to view the detailed status of DHCP snooping VLANs and ports:

```
S524DF4K15000024 # get switch dhcp-snooping database-summary

snoop-enabled-vlans             : 10
verifysrcmac-enabled-vlans      : 10
option82-enabled-vlans          : 10
option82-trust-enabled-intfs    :
trusted ports    :
untrusted ports  : port1 port2 port3 port4 port5 port6 port9 port10 port11 port12
               port13 port14 port15 port16 port17 port18 port19 port20 port21 port22
               port23 port24 port25 port26 port27 port28 port29 port30
Client Database          : 0 / 8000
Server Database          : 0 / 1024
Limit Database           : 0 / 256
```

An entry in the DHCP snooping binding database that contains an * after the IP address indicates a temporary or incomplete entry. For example:

```
08:00:27:13:16:51 2000 100.0.0.159* 10 4 port4
```

The DHCP server has not acknowledged this entry yet. If the DHCP server does not acknowledge the entry within 10 seconds, the entry is removed from the database. If the DHCP server does acknowledge the entry within 10 seconds, the entry will be considered “complete” (that is, no * after the IP address), and a proper expiration time is assigned to it.

Use the following command to view the details of the DHCP-snooping client database:

```
FS1D243Z14000027 # get switch dhcp-snooping client-db-details

       mac        vlan     ip     lease(sec) expiry(sec) interface hostname  domainname     vendor
00:01:00:00:00:01 100 xxx.x.x.xxx   86400       86398      port3
00:03:00:00:00:03 100 xxx.x.x.x     86400       86394      port5
00:03:00:00:00:04 100 xxx.x.x.x     86400       86394      port5
```

Use the following command to view the details of the DHCP-snooping server database:

```
FS1D243Z14000027 # get switch dhcp-snooping server-db-details
mac             vlan    ip     interface status first-seen (sec) last-seen (sec) ACK  NAC  OFFER OTHER
00:11:01:00:00:01 30 192.168.5.2  port6  trusted   1503357551         0           12    0      8    0
```

Removing an entry from the DHCP snooping binding database

You can remove an IP address from the DHCP snooping binding database by specifying the associated VLAN ID and MAC address:

```
execute dhcp-snooping expire-client <1-4095> <xx:xx:xx:xx:xx:xx>
```

For example:

```
execute dhcp-snooping expire-client 100 01:23:45:67:89:01
```

Dynamic ARP inspection

Dynamic ARP Inspection (DAI) prevents man-in-the-middle attacks and IP address spoofing by checking that packets from untrusted ports have valid IP-MAC-address binding. To use DAI, you must first enable the DHCP snooping feature and then enable DAI for each VLAN. See DHCP snooping on page 110.

This chapter covers the following topics:

Configuring DAI

Configuring DAI consists of the following steps:

1. Enable DAI for each VLAN. By default, it is disabled.
2. Enable DAI for the switch interface. By default, all interfaces are in an untrusted state. You must explicitly configure the trusted interfaces.

Enable DAI for each VLAN

```
config switch vlan
   edit <vlan-id>
      set arp-inspection {enable | disable}
   next
end
```

Enable DAI for the switch interface

```
config switch interface
   edit <interface-name>
      set arp-inspection-trust <untrusted | trusted>
   next
end
```

Checking ARP packets

Use the following command to see how many ARP packets have been dropped or forwarded:

```
#diagnose switch arp-inspection status

vlan 100           arp-request               arp-reply
-------------------------------------------------------------------------
received                 0                         0
forwarded                0                         0
dropped                  0                         0
```

IGMP snooping

FortiSwitch uses the information passed in IGMP messages to optimize the forwarding of multicast traffic.

IGMP snooping allows the FortiSwitch to passively listen to the Internet Group Management Protocol (IGMP) network traffic between hosts and routers. The switch uses this information to determine which ports are interested in receiving each multicast feed. FortiSwitch can reduce unnecessary multicast traffic on the LAN by pruning multicast traffic from links that do not contain a multicast listener.

Essentially, IGMP snooping is a layer-2 optimization for the layer-3 IGMP.

The current version of IGMP is version 3, and FortiSwitch is also compatible with IGMPv1 and IGMPv2.

Here is the basic IGMP snooping operation:

1. A host expresses interest in joining a multicast group. (Sends or responds to a join message).
2. FortiSwitch creates an entry in the layer-2 forwarding table (or adds the host's port to an existing entry). The switch creates one table entry per VLAN per multicast group.
3. FortiSwitch removes the entry when the last host leaves the group (or when the entry ages out).

In addition, you can configure FortiSwitch to send periodic queries from all ports in a specific VLAN to request IGMP reports. FortiSwitch uses the IGMP reports to update the layer-2 forwarding table.

This chapter covers the following topics:

Limitations

- You must enable the IGMP snooping function (using the `igmp-snooping enable` command) before you configure a multicast router port interface.
- Enabling the `set flood-unknown-multicast` command and then disabling it disrupts the forwarding of unknown multicast traffic to mRouter ports for a short period, depending on the query interval, because the mRouter ports need to be relearned.
- Currently, IGMPv3 (source-specific) is not fully supported. FortiSwitchOS can identify the IGMPv3 query/report messages, but the multicast group creation and traffic replication are based on the multicast group address and VLAN only (IGMPv2 operation).
- The IGMP snooping entries are added based on multicast group MAC address.
- Starting with release 3.5.2, the following snooping table limits apply:

Platform Series	IGMP Snooping Table Limit
124	1024
200	1024
400	1024
500	1024
1024 and 1048	4096
3032	8192

NOTE: Until FortiSwitch Release 3.5.1, the table limits were hardware only. The software limit for all platforms was 8192.

Configuring IGMP snooping

Configuring IGMP snooping consists of the following major steps:

1. Configure IGMP snooping on a global level.
2. Assign VLANs and enable IGMP snooping on the interfaces.
3. Configure IGMP snooping on the VLANs.

NOTE: IGMP snooping configured under "vlan enable" + "port based disable," does not work well; only "vlan level enable" + "port level enable" can make snooping work. So, because the port is "disabled" by default, you must enable IGMP snooping on both the VLAN and the port.

1. Configure IGMP snooping on a global level

By default, the maximum time (`aging-time`) that multicast snooping entries without any packets are kept is for 300 seconds. This value can be in the range of 15-3,600 seconds. By default, `flood-unknown-multicast` is disabled, and unregistered multicast packets are forwarded only to mRouter ports. If you enable `flood-unknown-multicast`, unregistered multicast packets are forwarded to all ports in the VLAN.

Using the CLI:

```
config switch igmp-snooping globals
   set aging-time <15-3600>
   set flood-unknown-multicast {enable | disable}
end
```

For example:

```
config switch igmp-snooping globals
   set aging-time 500
   set flood-unknown-multicast enable
end
```

2. Enable IGMP snooping on the interfaces

Enable IGMP snooping on a specified switch interface. The default is enabled.

Using the Web-based manager:

1. Go to **Switch > Interface > Interface** or **Switch > Interface > Trunk**.
2. Select an interface.
3. Click **Edit**.
4. Select the **IGMP Snooping** check box.
5. If needed, select the **Flood Reports** and **Flood Traffic** check boxes.
6. Click **OK**.

Using the CLI:

```
config switch interface
   edit <port>
      set native-vlan <vlan-id>
      set igmp-snooping {enable | disable}
      set igmps-flood-reports {enable | disable}
   next
end
```

For example:

```
config switch interface
   edit port10
      set native-vlan 30
      set igmp-snooping enable
   next
   edit port2
      set native-vlan 30
      set igmp-snooping enable
   next
   edit port4
      set native-vlan 30
      set igmp-snooping enable
   next
   edit port6
      set native-vlan 30
      set igmp-snooping enable
      next and
   edit port8
      set native-vlan 30
      set igmp-snooping enable
   next
end
```

Use the following command to clear the learned/configured multicast group from an interface:

```
execute clear switch igmp-snoop
```

3. Configure IGMP snooping on the VLANs

Enable IGMP snooping on a specified VLAN. The default is disabled.

You can define static groups for particular multicast addresses in a VLAN that has IGMP snooping enabled. The range of multicast addresses (mcast-addr) from 224.0.0.1 to 224.0.0.255 cannot be used. You can specify multiple ports in the static group, separated by a space. The trunk interface can also be included in a static group.

Using the CLI:

```
config switch vlan
   edit <vlan-id>
      set igmp-snooping {enable |disable}
      config igmp-static-group
         edit <group-name>
            set mcast-addr <multicast-address>
            set members <interface>
         next
      end
   next
end
```

For example, to configure two static groups for the same VLAN:

```
config switch vlan
   edit 30
      set igmp-snooping enable
      config igmp-static-group
         edit g239-1-1-1
            set mcast-addr 239.1.1.1
            set members port2 port5 port28
         next
         edit g239-2-2-2
            set mcast-addr 239.2.2.2
            set members port5 port10 trunk-1
         next
      end
   next
end
```

Check the IGMP snooping configuration

Use the following command to display information about IGMP snooping:

```
# get switch igmp-snooping (globals | group | interface | static-group)
```

- **globals**: display the IGMP snooping global configuration on the FortiSwitch
- **group**: display a list of learned groups
- **interface**: display the configured IGMP snooping interfaces and their current state
- **static-group**: display the list of configured static groups

Display the IGMP snooping global settings:

```
FS1D243Z13000023 # get switch igmp-snooping globals
aging-time : 300
flood-unknown-multicast: disabled
```

Display the learned multicast groups:

```
FS1D243Z13000023 # get switch igmp-snooping group
```

```
Number of Groups: 7
port of-port VLAN GROUP Age
(__port__9) 1 23 231.8.5.4 16
(__port__9) 1 23 231.8.5.5 16
(__port__9) 1 23 231.8.5.6 16
(__port__9) 1 23 231.8.5.7 16
(__port__9) 1 23 231.8.5.8 16
(__port__9) 1 23 231.8.5.9 16
(__port__9) 1 23 231.8.5.10 16
(__port__43) 3 23 querier 17
(__port__14) 8 --- flood-reports ---
(__port__10) 2 --- flood-traffic ---
```

Display the list of configured static groups:

```
FS1D243Z13000023 # get swi igm static-group

VLAN ID Group-Name     Multicast-addr  Member-interface
_______ _____________ _______________ ____________________________
11      g239-1        239:1:1:1       port6 trunk-2
11      g239-11       239:2:2:11      port26 port48 trunk-2
40      g239-1        239:1:1:1       port5 port25 trunk-2
40      g239-2        239:2:2:2       port25 port26
```

Configuring the IGMP querier

To use the IGMP querier, you need to configure how often IGMP queries are sent, enable the IGMP querier for a specific VLAN, and specify the address for the IGMP querier.

Use the following commands to specify how many seconds are between IGMP queries. The default is 120 seconds.

```
config switch igmp-snooping globals
   set query-interval <10-1200>
end
```

For example:

```
config switch igmp-snooping globals
   set aging-time 150
   set flood-unknown-multicast enable
   set query-interval 200
end
```

Use the following commands to enble the IGMP querier for a specific VLAN and specify the address that IGMP reports are sent to:

```
config switch vlan
   edit 100
      set igmp-snooping {enable | disable}
      set igmp-snooping-querier {enable | disable}
      set querier-addr <IPv4_address>
   next
end
```

For example:

```
config switch vlan
   edit 100
      set igmp-snooping enable
      set igmp-snooping-querier enable
      set querier-addr 1.2.3.4
   next
end
```

Configuring mRouter ports

Use the following commands to configure a FortiSwitch port as an mRouter port:

NOTE: These settings are not per-VLAN, so the port will act as a querier/mRouter port for all of its associated VLANs.

```
config switch interface
   edit <port>
      set igmp-snooping enable
      set igmps-flood-reports enable
      set igmps-flood-traffic enable
   next
end
```

Private VLANs

A private VLAN (PVLAN) divides the original VLAN (termed the primary VLAN) into sub-VLANs (secondary VLANs), while retaining the existing IP subnet and layer-3 configuration. Unlike a regular VLAN, which is a single broadcast domain, a PVLAN partitions one broadcast domain into multiple smaller broadcast subdomains.

After a PVLAN VLAN is configured, the primary VLAN forwards frames downstream to all secondary VLANs.

There are two main types of secondary VLANs:

- **Isolated**: Any switch ports associated with an isolated VLAN can reach the primary VLAN, but not any other secondary VLAN. In addition, hosts associated with the same isolated VLAN cannot reach each other. Only one isolated VLAN is allowed in one PVLAN domain.
- **Community**: Any switch ports associated with a common community VLAN can communicate with each other and with the primary VLAN but not with any other secondary VLAN. You might have multiple distinct community VLANs within one PVLAN domain.

There are mainly two types of ports in a PVLAN: promiscuous (P-Port) and host.

- **Promiscuous Port (P-Port)**: The switch port connects to a router, firewall, or other common gateway device. This port can communicate with anything else connected to the primary or any secondary VLAN. In other words, it is a type of a port that is allowed to send and receive frames from any other port on the VLAN.
- **Host Ports** further divides into two types – isolated port (I-Port) and community port (C-port).
  - **Isolated Port (I-Port)**: Connects to the regular host that resides on isolated VLAN. This port communicates only with P-Ports.
  - **Community Port (C-Port)**: Connects to the regular host that resides on community VLAN. This port communicates with P-Ports and ports on the same community VLAN.

This chapter covers the following topics:

Creating and enabling a PVLAN

Using the Web-based manager:

1. Go to **Switch** > **VLAN** > **PVLAN**.
2. Click **Create New** for a new PVLAN.
3. Enter the VLAN identifier.
4. Enter a description for the new PVLAN.
5. Select **Enable** to enable the new PVLAN.
6. Enter a single VLAN identifier for the isolated sub-VLAN.
7. If needed, enter one VLAN identifier or multiple VLAN identifiers for a common community sub-VLAN.
8. Click **OK**.

Configuring the PVLAN ports

Using the Web-based manager:

1. Go to **Switch** > **Interface** > **Interface**.
2. Select the port to configure.
3. Click **Edit**.
4. Select if the port is a promiscuous port or part of a sub-VLAN.
5. For a promiscuous port, select the primary VLAN identifier.
6. For a port that is part of a sub-VLAN, select the primary VLAN identifier and the sub-VLAN identifier.
7. Click **OK**.

Private VLAN example

1. Enabling a PVLAN:

```
config switch vlan
   edit 1000
      set private-vlan enable
      set isolated-vlan 101
      set community-vlans 200-210
   end
end
```

2. Configuring the PVLAN ports:

```
config switch interface
   edit "port2"
      set private-vlan promiscuous
      set primary-vlan 1000
   next
   edit "port3"
      set private-vlan sub-vlan
      set primary-vlan 1000
      set sub-vlan 200
   next
   edit "port7"
      set private-vlan sub-vlan
      set primary-vlan 1000
      set sub-vlan 101
   next
   edit "port19"
      set private-vlan promiscuous
      set primary-vlan 1000
   next
   edit "port20"
      set private-vlan sub-vlan
      set primary-vlan 1000
      set sub-vlan 101
   next
   edit "port21"
```

```
        set private-vlan sub-vlan
        set primary-vlan 1000
        set sub-vlan 101
    end
end
```

QoS settings

Quality of service (QoS) provides the ability to set particular priorities for different applications, users, or data flows.

QoS involves the following elements:

- **Classification** is the process of determining the priority of a packet. This can be as simple as trusting the QoS markings in the packet header when it is received and so accept the packet. Alternatively, it can hinge on criteria (such as incoming port, VLAN, or service) that are defined by the network administrator.
- **Marking** involves setting bits in the packet header to indicate the priority of this packet.
- **Queuing** involves defining priority queues to ensure that packets marked as high priority take precedence over those marked as lower priority. If network congestion becomes so severe that packet drops are inevitable, the queuing process will also select the packets to drop.

FortiSwitch supports the following QoS configuration capabilities:

- Mapping the IEEE 802.1p and layer-3 QoS values (Differentiated Services and IP Precedence) to an outbound QoS queue number.
- Providing eight egress queues on each port.
- Policing the maximum data rate of egress traffic on the interface.

This chapter covers the following topics:

Classification

The IEEE 802.1p standard defines a class of service (CoS) value (ranging from 0-7) that is included in the Ethernet frame. The Internet Protocol defines the layer-3 QoS values that are carried in the IP packet (Differentiated Services, IP Precedence). FortiSwitch provides configurable mappings from CoS or IP-DSCP values to egress queue values.

Fortinet recommends that you do not enable trust for both Dot1p and DSCP at the same time on the same interface. If you do want to trust both Dot1p and IP-DSCP, the switch uses the latter value (DSCP) to determine the queue. The switch will use the Dot1p value and mapping only if the packet contains no DSCP value. For details, refer to Determining the egress queue on page 127.

Marking

FortiSwitchOS supports two ways to indicate the priority of outgoing packets:

- **CoS marking**: The priority is set with the CoS value of the 802.1Q tag. The range of CoS values is 0-7.
- **Differential service code point (DSCP) marking**: The priority is set with the DSCP value in the IP header. The range of DSCP values is 0-63.

You can use one of these methods or both methods.

Whether the CoS or DSCP values of inbound packets are remarked is subject to the classification by ACL rules for the ingress interfaces. When CoS or DSCP marking take place, the outbound queuing is not impacted, meaning it is still based on trust maps and the original CoS or DSCP values, as described in Determining the egress queue on page 127.

The following example shows how to use the CLI to configure an ACL policy to mark the CoS and DSCP values of inbound packets to 4 and 48 on port1 when their CoS values are 2:

```
config switch acl policy
   edit 10
      config action
         set count enable
         set remark-cos 4
         set remark-dscp 48
      end
      config classifier
         set cos 2
      end
      set ingress-interface "port1"
   next
end
```

Queuing

Queuing determines how queued packets on an egress port are served. Each egress port supports eight queues, and three scheduling modes are available:

- **Strict Scheduling**: The queues are served in descending order (of queue number), so higher number queues receive higher priority. The purpose of the strict scheduling mode is to provide lower latency service to higher classes of traffic. However, if the interface experiences congestion, the lower priority traffic could be starved.
- **Simple Round Robin (RR)**: In round robin mode, the scheduler visits each backlogged queue, servicing a single packet from each queue before moving on to the next one. The purpose of round robin scheduling is to provide fair access to the egress port bandwidth.
- **Weighted Round Robin (WRR)**: Each of the eight egress queues is assigned a weight value ranging from 0 to 63. The purpose of weighted round robin scheduling is to provide prioritized access to the egress port bandwidth, such that queues with higher weight get more of the bandwidth, but lower priority traffic is not starved.

Determining the egress queue

To determine the egress queue value for the packet, FortiSwitch uses the configured trust values (and mappings) on the port and the QoS/CoS fields in the packet.

Packets with DSCP and CoS values

If the port is set to trust DSCP, the switch uses this value to find the queue assignment in the DSCP map for the port.

If the port is set to trust Dot1p and **not** to trust DSCP, the switch uses the packet's CoS value to look up the queue assignment in the Dot1p map for the port.

If the port is **not** set to trust Dot1p, the switch uses the default queue 0.

Packets with a CoS value but no DSCP value

The switch ignores the trust DSCP value.

- If the port is set to trust Dot1p, the switch uses the packet's CoS value to look up the queue assignment in the Dot1p map for the port.
- If the port is **not** set to trust Dot1p, the switch uses the default queue 0.

Packets with a DSCP value but no CoS value

If the port is set to trust DSCP, the switch uses the packet's DSCP value to look up the queue assignment in the DSCP map for the port.

If the port is set to trust Dot1p but **not** to trust DSCP, the switch uses the default CoS value of the port to look up the queue assignment in the Dot1p map for the port.

If the port is **not** set to trust Dot1p, the switch uses the default queue 0.

Configuring FortiSwitch QoS

This section provides procedures for the following configuration tasks:

- Configure a Dot1p map on page 128
- Configure a DSCP map on page 128
- Configure the egress QoS policy on page 129
- Configure the egress drop mode on page 130
- Configure the switch ports on page 131
- Configure QoS on trunks on page 131
- Configure QoS on VLANs on page 132
- Configure CoS and DSCP markings on page 132

Configure a Dot1p map

Using the Web-based manager:

1. Go to **Switch > QoS > 802.1p config**.
2. Click **Create New**.
3. Enter the name of your Dot1p map.
4. Enter a description of your Dot1p map.
5. Select the queue number for each priority.
6. Click **OK**.

Values that are not explicitly included in the map will follow the default mapping, which maps each priority (0-7) to queue 0. If an incoming packet contains no CoS value, the switch assigns a CoS value of zero.

Using the CLI:

To configure a Dot1p map, which defines a mapping between IEEE 802.1p CoS values (from incoming packets on a trusted interface) and the egress queue values, enter the following:

```
config switch qos dot1p-map
   edit <dot1p map name>
      set description <text>
      set [priority-0|priority-1|priority-2|....priority-7] <queue number>
   next
end
```

Example:

```
config switch qos dot1p-map
   edit "test1"
      set priority-0 queue-2
      set priority-1 queue-0
      set priority-2 queue-1
      set priority-3 queue-3
      set priority-4 queue-4
      set priority-5 queue-5
      set priority-6 queue-6
      set priority-7 queue-7
   next
end
```

Values that are not explicitly included in the map will follow the default mapping, which maps each priority (0-7) to queue 0. If an incoming packet contains no CoS value, the switch assigns a CoS value of zero.

Use the `set default-cos <port>` command to set a different default CoS value, ranging from 0 to 7:

```
config switch interface
   edit port1
      set default-cos <0-7>
```

Configure a DSCP map

A DSCP map defines a mapping between IP precedence or DSCP values and the egress queue values.

Using the Web-based manager:

1. Go to **Switch > QoS > IP precedence/DSCP**.
2. Click **Create New**.
3. Enter the name of your DCSP map.
4. Enter a description of your DCSP map.
5. Select which queue to configure.
6. Select the differentiated services to use.
7. Select the IP precedence to use.
8. Enter the raw values to use.
9. Click **OK**.

Using the CLI:

```
config switch qos ip-dscp-map
   edit <ip-dscp map name>
      set description <text>
      config map
      edit <entry-name1>
         set diffserv [ [ AF11 | AF12 | AF13 | AF21 | AF22 | AF23 | AF31 | AF32 | AF33 |
             AF41 | AF42 | AF43 | CS0 | CS1 | CS2 | CS3 | CS4 | CS5 | CS6 | CS7 | EF ]
         set ip-precedence [ Network Control | Internetwork Control | Critic/ECP | Flash
             Override | Flash, Immediate | Priority | Routine ]
         set value <dscp raw value>
         set cos-queue <queue number>
      next
   end
end
```

The following example defines a mapping for two of the DSCP values:

```
config switch qos ip-dscp-map
   edit "m1"
      config map
         edit "e1"
            set cos-queue 0
            set ip-precedence Immediate
         next
         edit "e2"
            set cos-queue 3
            set value 13
         next
      end
   next
end
```

Configure the egress QoS policy

In a QoS policy, you set the scheduling mode (Strict, Round Robin, or Weighted Round Robin) for the policy, and configure one or more CoS queues.

A valid set of values include the following:

- min-rate: minimum rate in kbps
- max-rate: maximum rate in kbps
- drop policy: taildrop or random early detection
- weight value (applicable if the policy schedule is weighted)

Using the Web-based manager:

1. Go to **Switch > QoS > Egress Policy**.
2. Click **Create New**.
3. Enter the name of your QoS egress policy.
4. Select the scheduling mode to use.
5. For each queue, enter a description and select the drop policy to use.
6. For each queue, enter the minimum rate in kbps, maximum rate in kbps, and weight value.
7. Click **OK**.

Using the CLI:

```
config switch qos qos-policy
   edit < policy.name >
      set schedule [ strict | round-robin | weighted ]
      config cos-queue
      edit [queue0 ... queue7]
         set description <text>
         set min-rate <rate kbps>
         set max-rate <rate kbps>
         set drop-policy {taildrop | random-early-detection | weighted-random-early-
             detction}
         set weight <value>
      end
   next
end
```

Configure the egress drop mode

NOTE: The egress-drop-mode command is available only for the 1024/1048/3032/5xx series.

When there are too many packets going through the same egress port, you can choose whether packets are dropped on ingress or egress.

Use the following commands to set the drop mode:

```
config switch physical-port
   edit <port>
      set egress-drop-mode <disabled | enabled>
   end
```

Variable	Description
disabled	Drop packets on ingress.
enabled	Drop packets on egress.

NOTE: Because too many packets are going through the same egress port, you might want to use the pause frame for flow control on the ingress side. To see the pause frame on ingress, enable the flow control “tx” on the ingress interface and disable egress-drop-mode on the egress interface.

Configure the switch ports

You can configure the following QoS settings on a switch port or a trunk:

- trust dot1p values on ingress traffic and the dot1p map to use
- trust ip-dscp values on ingress traffic and the ip-dscp map to use. (**NOTE:** Trust the dot1p values **or** the ip-dscp values but not both.)
- an egress policy for the interface
- a default CoS value (for packets with no CoS value)

If neither of the trust policies is configured on a port, the ingress traffic is mapped to queue 0 on the egress port.

If no egress policy is configured on a port, FortiSwitch applies the default scheduling mode (that is, round-robin).

Using the Web-based manager:

1. Go to **Switch > Interface > Interface**.
2. Select the switch port to update and click **Edit**.
3. Select the QoS egress policy.
4. Select the Dot1p map.
5. Select the DSCP map.
6. Click **OK**.

Using the CLI:

```
config switch interface
   edit <port>
      set trust-dot1p-map <map-name>
      set trust-ip-dscp-map <map-name>
      set qos-policy < policy-name >
      set default-cos <default cos value 0-7>
   next
end
```

Configure QoS on trunks

Configuring QoS on trunk interface follows the same configuration steps as for a switch port (configure a Dot1p/DSCP map and an egress policy).

When you add a port to a trunk, the port inherits the QoS configuration of the trunk interface. A port member reverts to the default QoS configuration when it is removed from the trunk interface.

Using the Web-based manager:

1. Go to **Switch > Interface > Trunk**.
2. Select the trunk to update and click **Edit**.
3. Select the QoS egress policy.
4. Select the Dot1p map.

5. Select the DSCP map.
6. Click **OK**.

Using the CLI:

The following example shows QoS configuration on a trunk interface:

```
config switch interface
   edit "tr1"
      set snmp-index 56
      set trust-dot1p-map "dot1p_map1"
      set default-cos 1
      set qos-policy "p1"
   next
end
```

When you configure an egress QoS policy with rate control on a trunk interface, that rate control value is applied to each port in the trunk interface. The FortiSwitch does not support an aggregate value for the whole trunk interface.

Configure QoS on VLANs

You can configure a CoS queue value for a VLAN by creating an ACL policy:

```
config switch acl policy
   edit 1
      config action
         set cos-queue 7
         set count enable
      end
      config classifier
         set vlan-id 200
      end
      set ingress-interface "port25"
end
```

Configure CoS and DSCP markings

You can classify a packet by matching the CoS value, DSCP value, or both CoS and DSCP values. You can also configure the action to set the CoS marking value, DSCP marking value, or both.

```
config switch acl policy
   edit <policy-id>
      config classifier
         set cos <802.1Q CoS value to match>
         set dscp <DSCP value to match>
      end
      config action
         set remark-cos <0-7>
         set remark-dscp <0-63>
      end
```

For example:

```
config switch acl policy
   edit 1
      config classifier
```

```
            set src-mac 11:22:33:44:55:66
            set cos 2
            set dscp 10
         end
         config action
            set count enable
            set remark-cos 4
            set remark-dscp 20
         end
      set ingress-interface port2
   end
```

Checking the QoS statistics

To check the statistics for all QoS queues, use the following command:

```
diagnose switch physical-ports qos-stats list
```

To check the statistics for QoS queues for specific ports, use the following command:

```
diagnose switch physical-ports qos-stats list <list_of_ports>
```

The output differs depending on the FortiSwitch model.

For example, for the 1xxxD, 3xxxD, and 5xxD FortiSwitch models:

```
diagnose switch physical-ports qos-stats list 1,3,4-6
```

```
port1 QoS Stats:

queue |      unicast pkts |     unicast bytes |    multicast pkts |   multicast
bytes
-------------------------------------------------------------------------------
0 |                  0 |                  0 |                  0 |                  0
1 |                  0 |                  0 |                  0 |                  0
2 |                  0 |                  0 |                  0 |                  0
3 |                  0 |                  0 |                  0 |                  0
4 |                  0 |                  0 |                  0 |                  0
5 |                  0 |                  0 |                  0 |                  0
6 |                  0 |                  0 |                  0 |                  0
7 |                  0 |                  0 |                  0 |                  0

queue |   ucast drop pkts |  ucast drop bytes |   mcast drop pkts |  mcast drop
bytes
-------------------------------------------------------------------------------
0 |                  0 |                  0 |                  0 |                  0
1 |                  0 |                  0 |                  0 |                  0
2 |                  0 |                  0 |                  0 |                  0
3 |                  0 |                  0 |                  0 |                  0
4 |                  0 |                  0 |                  0 |                  0
5 |                  0 |                  0 |                  0 |                  0
6 |                  0 |                  0 |                  0 |                  0
7 |                  0 |                  0 |                  0 |                  0
```

```
--------------------------------------------------------------------------------

port3 QoS Stats:

queue |      unicast pkts |     unicast bytes |    multicast pkts |   multicast
bytes
--------------------------------------------------------------------------------
0 |                 0 |                 0 |                 0 |                 0
1 |                 0 |                 0 |                 0 |                 0
2 |                 0 |                 0 |                 0 |                 0
3 |                 0 |                 0 |                 0 |                 0
4 |                 0 |                 0 |                 0 |                 0
5 |                 0 |                 0 |                 0 |                 0
6 |                 0 |                 0 |                 0 |                 0
7 |                 0 |                 0 |                 0 |                 0

queue |   ucast drop pkts |  ucast drop bytes |   mcast drop pkts |  mcast drop
bytes
--------------------------------------------------------------------------------
0 |                 0 |                 0 |                 0 |                 0
1 |                 0 |                 0 |                 0 |                 0
2 |                 0 |                 0 |                 0 |                 0
3 |                 0 |                 0 |                 0 |                 0
4 |                 0 |                 0 |                 0 |                 0
5 |                 0 |                 0 |                 0 |                 0
6 |                 0 |                 0 |                 0 |                 0
7 |                 0 |                 0 |                 0 |                 0
--------------------------------------------------------------------------------

port4 QoS Stats:

queue |      unicast pkts |     unicast bytes |    multicast pkts |   multicast
bytes
--------------------------------------------------------------------------------
0 |                 0 |                 0 |                 0 |                 0
1 |                 0 |                 0 |                 0 |                 0
2 |                 0 |                 0 |                 0 |                 0
3 |                 0 |                 0 |                 0 |                 0
4 |                 0 |                 0 |                 0 |                 0
5 |                 0 |                 0 |                 0 |                 0
6 |                 0 |                 0 |                 0 |                 0
7 |                 0 |                 0 |                 0 |                 0

queue |   ucast drop pkts |  ucast drop bytes |   mcast drop pkts |  mcast drop
bytes
--------------------------------------------------------------------------------
0 |                 0 |                 0 |                 0 |                 0
1 |                 0 |                 0 |                 0 |                 0
2 |                 0 |                 0 |                 0 |                 0
3 |                 0 |                 0 |                 0 |                 0
4 |                 0 |                 0 |                 0 |                 0
5 |                 0 |                 0 |                 0 |                 0
```

```
6 |                 0 |                 0 |                 0 |                 0
7 |                 0 |                 0 |                 0 |                 0
--------------------------------------------------------------------------------

port5 QoS Stats:

queue |      unicast pkts |     unicast bytes |    multicast pkts |   multicast
bytes
--------------------------------------------------------------------------------
0 |                 0 |                 0 |                 0 |                 0
1 |                 0 |                 0 |                 0 |                 0
2 |                 0 |                 0 |                 0 |                 0
3 |                 0 |                 0 |                 0 |                 0
4 |                 0 |                 0 |                 0 |                 0
5 |                 0 |                 0 |                 0 |                 0
6 |                 0 |                 0 |                 0 |                 0
7 |                 0 |                 0 |                 0 |                 0

queue |   ucast drop pkts |  ucast drop bytes |   mcast drop pkts |  mcast drop
bytes
--------------------------------------------------------------------------------
0 |                 0 |                 0 |                 0 |                 0
1 |                 0 |                 0 |                 0 |                 0
2 |                 0 |                 0 |                 0 |                 0
3 |                 0 |                 0 |                 0 |                 0
4 |                 0 |                 0 |                 0 |                 0
5 |                 0 |                 0 |                 0 |                 0
6 |                 0 |                 0 |                 0 |                 0
7 |                 0 |                 0 |                 0 |                 0
--------------------------------------------------------------------------------

port6 QoS Stats:

queue |      unicast pkts |     unicast bytes |    multicast pkts |   multicast
bytes
--------------------------------------------------------------------------------
0 |                 0 |                 0 |                 0 |                 0
1 |                 0 |                 0 |                 0 |                 0
2 |                 0 |                 0 |                 0 |                 0
3 |                 0 |                 0 |                 0 |                 0
4 |                 0 |                 0 |                 0 |                 0
5 |                 0 |                 0 |                 0 |                 0
6 |                 0 |                 0 |                 0 |                 0
7 |                 0 |                 0 |                 0 |                 0

queue |   ucast drop pkts |  ucast drop bytes |   mcast drop pkts |  mcast drop
bytes
--------------------------------------------------------------------------------
0 |                 0 |                 0 |                 0 |                 0
1 |                 0 |                 0 |                 0 |                 0
2 |                 0 |                 0 |                 0 |                 0
3 |                 0 |                 0 |                 0 |                 0
```

```
4 |                 0 |                 0 |                 0 |                 0
5 |                 0 |                 0 |                 0 |                 0
6 |                 0 |                 0 |                 0 |                 0
7 |                 0 |                 0 |                 0 |                 0
-------------------------------------------------------------------------------
```

For example, for the 4xxD, 4xxD-POE, 4xxD-FPOE, 2xxD, 2xxD-POE, and 2xxD-FPOE FortiSwitch models:

```
diagnose switch physical-ports qos-stats list 1,6,48

port1 QoS Stats:

queue |              pkts |             bytes |         drop pkts
-----------------------------------------------------------------
0 |              1073 |           1017488 |                 0
1 |                 0 |                 0 |                 0
2 |                 0 |                 0 |                 0
3 |                 0 |                 0 |                 0
4 |                 0 |                 0 |                 0
5 |                 0 |                 0 |                 0
6 |                 0 |                 0 |                 0
7 |             60678 |           7700394 |                 0
-----------------------------------------------------------------

port6 QoS Stats:

queue |              pkts |             bytes |         drop pkts
-----------------------------------------------------------------
0 |            104779 |          36489164 |                 0
1 |                 0 |                 0 |                 0
2 |               315 |            317960 |                 0
3 |                 0 |                 0 |                 0
4 |                 0 |                 0 |                 0
5 |            267121 |         267121000 |                 0
6 |            766734 |         766734000 |                 0
7 |            572592 |          42493451 |                 0
-----------------------------------------------------------------

port48 QoS Stats:

queue |              pkts |             bytes |         drop pkts
-----------------------------------------------------------------
0 |           1628754 |         131562453 |                 0
1 |                 0 |                 0 |                 0
2 |               400 |            400400 |                 0
3 |                 0 |                 0 |                 0
4 |           2054967 |        2054967000 |                 0
5 |            438759 |         438759000 |                 0
6 |            137577 |         137577000 |                 0
7 |            166364 |          72177282 |                 0
-----------------------------------------------------------------
```

Clearing the QoS statistics

The `diagnose switch physical-ports qos-stats clear` command is supported only for the 1xxxD, 3xxxD, and 5xxD FortiSwitch models. The `diagnose switch physical-ports qos-stats clear` command is not available for the 4xxD, 4xxD-POE, 4xxD-FPOE, 2xxD, 2xxD-POE, or 2xxD-FPOE FortiSwitch models.

To clear the statistics for the QoS queues for all ports, use the following command:

```
diagnose switch physical-ports qos-stats clear
```

To clear the statistics for the QoS queues for specified ports, use the following command:

```
diagnose switch physical-ports qos-stats clear <list_of_ports>
```

For example:

```
diagnose switch physical-ports qos-stats clear 1,3,4-6
```

sFlow

sFlow is a method of monitoring the traffic on your network to identify areas on the network that may impact performance and throughput. With sFlow you can export truncated packets and interface counters. FortiSwitch implements sFlow version 5 and supports trunks and VLANs.

This chapter covers the following topics:

About sFlow

sFlow uses packet sampling to monitor network traffic. The sFlow agent captures packet information at defined intervals and sends them to an sFlow collector for analysis, providing real-time data analysis. To minimize the impact on network throughput, the information sent is only a sampling of the data.

The sFlow collector is a central server running software that analyzes and reports on network traffic. The sampled packets and counter information, referred to as flow samples and counter samples, respectively, are sent as sFlow datagrams to a collector. Upon receiving the datagrams, the sFlow collector provides real-time analysis and graphing to indicate the source of potential traffic issues. sFlow collector software is available from a number of third-party software vendors.

Configuring sFlow

Configuration consists of the following steps:

1. Enable the sFlow agent.
2. Configure sampling information on the interfaces.

Configure sFlow agents

Use the following commands to configure an sFlow agent:

1. Set the IP address of the collector.
2. Set the collector port number, which is the destination port number in sFlow UDP packets. The default value is 6343.

Using the Web-based manager:

1. Go to **System > Network > sFlow**.
2. Set the collector IP address and port number.
3. Click **OK** to save the changes.

Using the CLI:

```
config system sflow
   set collector-ip <ip/hostname>
   set collector-port <port>
```

Configure the interfaces

Use the following commands to configure sFlow on a port:

- Enable sFlow on the port (by default, sFlow is disabled).
- Set the sample rate. An average of one out of `count` packets is randomly sampled. The rate ranges from 0-99999; the default is 512.
- Set the direction for capturing the traffic. sFlow can capture the ingress traffic (RX), the egress traffic (TX), or both (the default).
- Set the polling interval, which defines how often the switch sends interface counters to the collector. The range of values is 1-255 and default is 30.

Using the Web-based manager:

1. Go to **Switch > Interface > Interface.**
2. Select one or more ports to update and click **Edit**.
3. If you selected more than one port, the port names are displayed in the name field, separated by commas.
4. Set **Enable sFlow**.
5. Enter new values as required for Sample Rate, Sample Direction, and Polling Interval.
6. Click **OK** to save the changes.

Using the CLI:

```
config switch interface
   edit <port>
      set sflow-sampler [enabled | disabled]
      set sample-rate <count>
      set sample-direction [rx | tx | both]
      set polling-interval <interval>
```

NOTE: Ensure that you can use the exec command `ping collector_ip_address` to ping the collector from the FortiSwitch. Then, use the built-in sniffer to trace sFlow packets (`diag sniff packet <vlan_interface_name> "udp port 6343"`).

Checking the sFlow configuration

Use the following command to display the sFlow configuration:

```
get system sflow
```

Feature licensing

Advanced features (such as dynamic routing protocols) require a feature license.

This chapter covers the following topics:

About licenses

Each feature license is tied to the serial number of the FortiSwitch. Therefore, a feature license is valid on one system.

Configuring licenses

Configuration consists of the following steps:

1. Check license status.
2. Add a license.

Checking the license status

Using the Web-based manager:

1. Go to **System > Dashboard > Status**.
2. Check which licenses that are currently active.

Using the CLI:

```
execute license status
```

Adding a license

NOTE: Adding license keys causes the system to log you out.

Using the Web-based manager:

1. Go to **System > Config > License**.
2. Click **Add**.
3. Enter your license key.
4. Click **OK**.

Using the CLI:

```
execute license add <key>
```

Removing a license

Using the Web-based manager:

1. Go to **System > Config > License**.
2. Select a license to remove
3. Click **Delete**.
4. Click **OK** to acknowledge the warning.

NOTE: Deleting license keys causes the system to log you out before rebooting. You will lose all configurations related to the license.

Using the CLI:

```
execute license type <type> clear
```

Layer-3 interfaces

Fortinet data center switches support loopback interfaces and switched virtual interfaces (SVIs), both of which are described in this chapter.

This chapter covers the following topics:

Loopback interfaces

A loopback interface is a special virtual interface created in software that is not associated with any hardware interface.

Dynamic routing protocols typically use a loopback interface as a reliable IP interface for routing updates. You can assign the loopback IP address to the router rather than the IP address of a specific hardware interface. Services (such as Telnet) can access the router using the loopback IP address, which remains available independent of hardware interfaces status.

No limit exists on the number of loopback interfaces you can create.

A loopback interface does not have an internal VLAN ID or a MAC addresses and always uses a /32 network mask.

Configuring loopback interfaces

Using the Web-based manager:

1. Go to **System > Network > Interface** and select **Create New**.

Using the CLI:

```
config system interface
   edit "loopback"
      set ip 172.168.20.1 255.255.255.255
      set allowaccess ping https http ssh telnet
      set type loopback
      set snmp-index 28
   next
end
```

Switched virtual interfaces

A switched virtual interface (or SVI) is a logical interface that is associated with a VLAN and supports routing and switching protocols.

You can assign an IP address to the SVI to enable routing between VLANs. For example, SVIs can route between two different VLANs connected to a switch (no need to connect through a layer-3 router).

Configuring a switched virtual interface

Using the Web-based manager:

1. Go to **System > Network > Interface** and select **Create New**.
2. Provide the interface an appropriate name.
3. Set **Interface** to **internal**.
4. Set a **VLAN ID**.
5. Assign an **IP/Netmask**.
6. Set **Administrative Access** to allow ping, SSH, and Telnet.
7. Select **OK**.

Using the CLI:

Create a system interface. Give it an IP subnet and an associated VLAN:

```
config system interface
   edit <system interface name>
      set ip <IP address and mask>
      set vlanid <vlan>
      set allowaccess ping ssh telnet
```

Example SVI configuration

The following is an example CLI configuration for SVI static routing.

In this configuration, Server-1 is connected to switch Port1, and Server-2 is connected to switch Port2. Port1 is a member of VLAN 4000, and Port2 is a member of VLAN 2. Port1 is the gateway for Server-1, and port2 is the gateway for Server-2.

NOTE: For simplicity, assume that both port1 and port are on same switch.

1. Configure the native VLANs for Port 1 and Port 2:

```
config switch interface
   edit port1
      set native-vlan 4000
   edit port2
      set native-vlan 2
   end
```

2. Create L3 system interfaces that correspond to Port 1 (VLAN 4000) and Port 2 (VLAN 2):

```
config system interface
   edit vlan4000
```

```
      set ip 192.168.11.1/24
      set vlanid 4000
      set allowaccess ping ssh telnet
   next
      edit vlan2
         set ip 192.168.10.1/24
         set vlanid 2
         set allowaccess ping ssh telnet
   end
```

Viewing the SVI configuration

Display the status of SVI configuration using following command:

```
show system interface [ <system interface name> ]
```

Layer-3 routing in hardware

In Release 3.3.0 and later, some FortiSwitch models support hardware-based layer-3 forwarding.

For FortiSwitch models that support Equal Cost Multi-Path (ECMP) (see Feature matrix: FortiSwitchOS 3.6 on page 12), forwarding for all ECMP routes is performed in hardware.

For switch models that support hardware-based layer-3 forwarding but do not support ECMP, only one route to each destination will be hardware-forwarded. If you configure multiple routes to the same destination, you can configure a priority value for each route. Only the route with highest priority will be forwarded by the hardware. If no priority values are assigned to the routes, the most recently configured route is forwarded by the hardware.

Router activity

Logging allows you to review all router activity.

NOTE: Router logs are available only on supported platforms if you have the advanced features license.

To enable router logging:

1. Go to **Log > Log Config > Log Setting**.
2. Select **Event Logging**.
3. Select **Router activity event**.
4. Select **Apply**.

To view router logs:

1. Go to **Log > Log Config > Event Log > Router**.
2. Click **Download Raw Log** if you want to review the entries offline.

Equal cost multi-path (ECMP) routing

ECMP is a forwarding mechanism that enables load-sharing of traffic to multiple paths of equal cost. An ECMP set is formed when the routing table contains multiple next-hop address for the same destination with equal cost. Routes of equal cost have the same preference and metric value. If there is an ECMP set for an active route, the switch uses a hash algorithm to choose one of the next-hop addresses. As input to the hash, the switch uses one or more of the following fields in the packet to be routed:

- Source IP
- Destination IP
- Input port

Configuring ECMP

The switch automatically uses ECMP to choose between equal-cost routes.

This configuration value is system-wide. The source IP address is the default value.

Notes and Restrictions

When you configure a static route with a gateway, the gateway must be in the same IP subnet as the device. Also, the destination subnet cannot match any of device IP subnets in the switch.

When you configure a static route without a gateway, the destination subnet must be in the same IP subnet as the device.

Using the CLI:

```
config system settings
   set v4-ecmp-mode [ source-ip-based ] [ dst-ip-based ] [ port-based ]
end
```

Example ECMP configuration

The following is an example CLI configuration for ECMP forwarding.

In this configuration, ports 2 and 6 are routed ports. Interfaces I-RED and I-GREEN are routed VLAN interfaces. The remaining ports in the switch are normal layer-2 ports.

1. Configure native VLANs for ports 2, 6, and 9. Also configure the “internal” interface to allow native VLANs for ports 2, 6, and 9:

```
config switch interface
   edit port2
      set native-vlan 10
   edit port6
      set native-vlan 20
   edit port9
      set native-vlan 30
   edit internal
      set allowed-vlans 10,20,30
   end
```

2. Configure the system interfaces:

```
config system interface
   edit "internal"
      set type physical
   next
      edit "i-blue"
         set ip 1.1.1.1 255.255.255.0
         set allowaccess ping https http ssh snmp telnet
         set vlanid 10
         set interface internal
   next
      edit "i-red"
         set ip 172.16.11.1 255.255.255.0
         set allowaccess ping ssh telnet
         set vlanid 20
         set interface internal
   next
      edit "i-green"
         set ip 172.168.13.1 255.255.255.0
         set allowaccess ping https http ssh snmp telnet
         set vlanid 30
         set interface internal
   next
end
```

3. Configure static routes. This code configures multiple next-hop gateways for the same network:

```
config router static
   edit 1
      set device "mgmt"
      set gateway 10.105.0.1
   next
   edit 2
      set device "i-red"
      set dst 8.8.8.0/24
      set gateway 172.16.11.2
   next
   edit 3
      set device "i-green"
      set dst 8.8.8.0/24
      set gateway 172.168.13.2
   next
```

Viewing ECMP configuration

Display the status of the ECMP configuration using following command:

```
show system interface [ <system interface name> ]
```

Bidirectional forwarding detection

FortiSwitchOS v3.4.2 and later supports static bidirectional forwarding detection (BFD), a point-to-point protocol to detect faults in the datapath between the endpoints of an IETF-defined tunnel (such as IP, IP-in-IP, GRE, and MPLS LSP/PW).

BFD defines demand mode and asynchronous mode operation. The FortiSwitch supports asynchronous mode. In this mode, the systems periodically send BFD control packets to one another, and if a number of those packets in a row are not received by the other system, the session is declared to be down.

BFD packets are transported using UDP/IP encapsulation and BFD control packets are identified using well-known UDP destination port 3784 (**NOTE:** BFD echo packets are identified using 3785).

BFD packets are not visible to the intermediate nodes and are generated and processed by the tunnel end systems only.

Configuring BFD

Use the following steps to configure BFD:

1. Configure the following values in the system interface:
   - **Enable BFD**: Set to **enable** or set to **global** to inherit the global configuration value.
   - **Desired min TX interval**: This is the minimum interval that the local system would like to use between transmission of BFD control packets. Value range is 200 ms – 30,000 ms. Default value is 250.
   - **Required min RX interval**: This is the minimum interval that the local system can support between receipt of BFD control packets. If you set this value to zero, the remote system will not transmit BFD control packets. The value range is 200 ms – 30000 ms. The default value is 250.
   - **Detect multi**: This is the detection time multiplier. The negotiated transmit interval multiplied by this value is the Detection Time for the receiving system. The value range is 1 – 20. The default is 3.
2. Enable BFD in the static router configuration.

Using the CLI:

```
config system interface
   edit <system interface name>
      set bfd [enable| disable | global]
      set bfd-desired-min-tx <number of ms>
      set bfd-required-min-rx <number of ms>
      set bfd-detect-multi [1…20]
next
config router static
edit 1
   set bfd enable
```

Viewing BFD configuration

Display the status of BFD sessions using following command:

```
get router info bfd neighbor [ <IP address of neighbor>]
```

```
OurAddr          NeighAddr        LD/RD    State    Int
192.168.15.2     192.168.15.1     1/4      UP       vlan2000
192.168.16.2     192.168.16.1     2/2      UP       vlan2001
```

Use the following command to display additional details:

```
get router info bfd neighbor detail
```

IP-MAC binding

Use IP-MAC binding to prevent ARP spoofing.

The port accepts a packet only if the source IP address and source MAC address in the packet match an entry in the IP-MAC binding table.

You can enable/disable IP-MAC binding for the whole switch, and you can override this global setting for each port.

Configuring IP-MAC binding

Use the following steps to configure IP-MAC binding:

1. Enable the IP-MAC binding global setting.
2. Create the IP-MAC bindings. You can activate each binding individually.
3. Set each port to follow the global setting. You can also override the global setting for individual ports by enabling or disabling IP-MAC binding for the port.

Using the Web-based manager:

Enable the IP-MAC binding global setting:

1. Go to **Switch** > **IP MAC Binding** > **Settings**.
2. Click **Enable** to enable IP-MAC binding.
3. Click **Apply** to save the change.

Create the IP-MAC bindings:

1. Go to **Switch** > **IP MAC Binding** > **Bindings**.
2. Click **Create New** to create a new binding.

Enable IP-MAC binding on the interface:

1. Go to **Switch** > **Interfaces** > **Interface**.
2. Edit the interface to be configured.
3. Select one of the IP-MAC binding settings.

Using the CLI:

```
config switch global
   set ip-mac-binding [enable| disable]

config switch ip-mac-binding
   edit 1
      set ip <IP address and network mask>
      set mac <MAC address>
```

```
      set status (enable| disable)
   next
end
config switch interface
   edit <port>
      set ip-mac-binding (enable| disable | global)
   edit <trunk name>
      set ip-mac-binding (enable| disable | global)
```

Notes

For a switch port, the default IP-MAC binding value is disabled.

When you configure a trunk, the trunk follows the global value by default. You can also explicitly enable or disable IP-MAC binding for a trunk, as shown in the CLI configuration.

When you add member ports to the trunk, all ports take on the trunk setting. If you later remove a port from the trunk group, the port is reset to the default value (disabled).

No duplicate entries are allowed in the mapping table.

Rules are disabled by default. You need to explicitly enable each rule.

The mapping table holds up to 1024 rules.

Viewing IP-MAC binding configuration

Display the status of IP-MAC binding using the following command:

```
show switch ip-mac-binding <entry number>
```

DHCP relay

DHCP clients send broadcast requests to a DHCP server. Without DHCP relay, the DHCP client and server must be on the same subnet. DHCP relay behaves as a proxy between DHCP clients and a DHCP server on a different subnet.

When the DHCP relay receives a DHCP request from a host on an inside interface, it forwards the request to one of the specified DHCP servers on an outside interface. When the DHCP server responds to the client request, the DHCP relay forwards the response back to DHCP client.

This chapter covers the following topics:

Detailed operation

DHCP relay operates as follows:

1. DHCP client C broadcasts a DHCP/BOOTP discover message on its subnet.
2. The relay agent examines the gateway IP address field in the DHCP/BOOTP message header. If the field has an IP address of 0.0.0.0, the agent fills it with the relay agent's or router's IP address and forwards the message to the remote subnet of the DHCP server.
3. When DHCP server receives the message, it examines the gateway IP address field for a DHCP scope that can be used by the DHCP server to supply an IP address lease.
4. If DHCP server has multiple DHCP scopes, the address in the gateway IP address field (GIADDR) identifies the DHCP scope from which to offer an IP address lease.
5. DHCP server sends an IP address lease offer (DHCPOFFER) directly to the relay agent identified in the gateway IP address (GIADDR) field.
6. The router then relays the address lease offer (DHCPOFFER) to the DHCP client.

Notes

DHCP relay service supports up to 8 relay targets per interface.

Each target is sent a copy of the DHCP message.

Configuring DHCP relay

You can configure DHCP relay on any layer-3 interface.

Using the Web-based manager:

1. Go to **System > Network > Interface**.
2. Select an interface.
3. Click **Edit**.
4. Select the **DHCP Relay** check box.
5. Enter the IP addresses for the relay servers, separated by a space.
6. If you want to include Option-82 data, select the **Option-82** check box.
7. Click **OK**.

Using the CLI:

```
config system interface
   edit <interface-name>
      set dhcp-relay-service (enable | disable)
      set dhcp-relay-ip <ip-address1> [<ip-address2> ... <ip-address8>]
      set dhcp-relay-option82 (enable | disable)
   next
end
```

Configuration example

In the following example, the DHCP server has address 192.168.23.2:

```
edit "v15-p15"
  set dhcp-relay-service enable
  set dhcp-relay-ip "192.168.23.2"    -> the DHCP server address
  set ip 192.168.15.1 255.255.255.0   -> the DHCP client subnet
  set allowaccess ping ssh snmp telnet
  set snmp-index 53
  set vlanid 15
  set interface "internal"
end
```

OSPF routing

NOTE: You must have an advanced features license to use OSPF routing.

Open shortest path first (OSPF) is a link-state interior routing protocol that is widely used in large enterprise organizations. OSPF provides routing within a single autonomous system (AS). This differs from BGP, which provides routing between autonomous systems.

An OSPF AS can contain only one area, or it may consist of a group of areas connected to a backbone area. A router connected to more than one area is an area border router (ABR). Routing information is contained in a link state database. Routing information is communicated between routers using link state advertisements (LSAs).

The main benefit of OSPF is that it detects link failures in the network quickly and converges network traffic successfully within seconds without any network loops. Also, OSPF has features to control which routes are propagated to contain the size of the routing tables.

You can enable bidirectional forwarding detection (BFD) with OSPF. BFD is used to quickly locate hardware failures in the network. Routers running BFD communicate with each other, and, if a timer runs out on a connection, that router is declared to be down. BFD then communicates this information to OSPF, and the routing information is updated.

FortiSwitch supports the following capabilities:

- Support for OSPFv2.
- OSPF neighbors support authentication
- Supports NSSA
- Supports BFD
- Supports stub area
- Network scaling

NOTE: OSPF MIBs are not supported in this release.

For additional information about OSPF routing, see the OSPF section of the FortiOS Handbook.

This chapter covers the following topics:

Terminology

Link State: Information is shared between directly connected routers. This information propagates throughout the network unchanged and is also used to create a shortest path first (SPF) tree.

Autonomous System (AS) : A network under a common network administration.

Area: You can divide a large network into areas to limit the number of link-state updates.

Cost: The routing metric used by OSPF. Lower costs are always preferred. You can configure the cost or use the interface default.

Router ID: Each OSPF router requires a unique router ID. For FortiSwitch, the unique router ID must be assigned manually.

Adjacency: When two OSPF routers have exchanged information and have the same topology table.

Topology Table: Also called the link-state table. This table contains information about every link in the network. The SPF algorithm uses the link-state information to calculate the best route to each destination.

Designated Router (DR): This router is responsible for ensuring adjacencies between all neighbors on a multi-access network (such as Ethernet). This ensures all routers do not need to maintain full adjacencies with each other. The DR is selected based on the router priority. In a tie, the router with the highest router ID is selected.

Backup DR: A backup router designed to perform the same functions in case the DR fails.

Link-State Advertisement (LSA): The method used by each router to share its routing topology with other routers in the same area.

Area Border Router (ABR): Router located on the border of one or more OSPF areas that connects those areas to the backbone area.

AS Boundary Router (AS BR): ABR located between an OSPF autonomous system and a non-OSPF network.

How OSPF works

Areas

An OSPF implementation consists of one or more areas. An area consists of a group of contiguous networks. If you configure more than one area, Area Zero is always the backbone area. An ABR links one or more areas to the OSPF backbone area.

FortiSwitch supports different types of areas—stub areas, Not So Stubby areas (NSSA), and regular areas. A stub area is an interface without a default route configured. NSSA is a type of stub area that can import AS external routes and send them to the backbone but cannot receive AS external routes from the backbone or other areas. All other areas are considered regular areas.

Adjacencies

When an OSPF router boots up, it sends OSPF Hello packets to find neighbors on the same network. Neighbors exchange information, and the Link State databases of both neighbors are synchronized. At this point, these neighbors are said to be adjacent.

For two OSPF routers to become neighbors, the following conditions must be met:

- The subnet number and subnet mask for the interface must match in both routers.
- The Hello interval and Dead interval values must match.
- The routers must have the same OSPF area ID.
- If authentication is used, they must pass authentication checks.

In OSPF, routing protocol packets are only passed between adjacent routers.

Configuring OSPF

Using the Web-based manager:

1. Create a switched virtual interface. See Configuring a switched virtual interface on page 143.
2. Go to **Router > Router > OSPF**.
3. Enter a unique 32-bit number in dotted decimal format for the router identifier. **NOTE:** Without a router identifier, OSPF routing will not work.
4. Select an area and click **Create New**.
   - Select if the area is a stub area, NSSA, or a regular area.
   - If you want routing authentication, select **MD5** or **Text**.
   - Click **OK**.
5. Under **Networks**, click **Create New**.
   - Enter the IP address and netmask, separated with a space. Use an IP address that includes the switched virtual interface.
   - Select the area that you created.
   - Click **OK**.
6. Under **Interfaces**, click **Create New**.
   - Enter a descriptive name for the OSPF interface name.
   - Select the switched virtual interface that you created.
   - Select the same type of authentication that you selected for the area.
   - If you want static bidirectional forwarding detection, select **Enable** or **Global**.
   - Enter the maximum transmission unit.
   - Enter the cost.
   - Enter the number of seconds between Hello packets being sent.
   - Enter the number of seconds that a Hello packet is not received before the OSPF router decides that a neighbor has failed.
   - Click **OK**.
7. Click **Advanced Options (BFD, Default Route, Redistribution)**.
   - If you are going to advertise non-OSPF routes within OSPF, enter the metric (cost) for other routing protocols.
   - Click **Apply**.

Using the CLI:

Configuring OSPF on FortiSwitch includes the following major steps:

1. Enter the OSPF configuration mode.
2. Set the router identifier. Each router must have a unique 32-bit number. **NOTE:** Without a router identifier, OSPF routing will not work.
3. Create an area. You must create at least one area.
4. Configure the network. Attach one or more networks to each area.
5. Configure an interface to a peer OSPF router.
6. Redistribute non-OSPF routes. Advertise these non-OSPF routes within OSPF.

1. Enter the OSPF configuration mode

Enter the OSPF configuration mode to access all of the OSPF configuration commands:

```
# config router ospf
```

2. Set the router identifier

Each router within an area must have a unique 32-bit number. The router identifier is written in dotted decimal format, but it is not an IPv4 address. **NOTE:** Without a router identifier, OSPF routing will not work.

```
set router-id <router-id>
```

For example:

```
# config router ospf
(ospf) # set router-id 1.1.1.2
```

3. Create an area

You must create at least one area. The area number is written in dotted decimal format (for example, configure area 100 as 0.0.0.100).

```
config area
   edit <area number>
      set authentication {md5 | none | text}
      set shortcut (default | disable | enable)
      set type {nssa | regular | stub}
end
```

For example:

```
(ospf) # config area
(area) # edit 0.0.0.4
(0.0.0.4) # set type nssa
(0.0.0.4) # set authentication md5
```

4. Configure the network

Use this subcommand to identify the OSPF-enabled interfaces. The prefix length in the interface must be equal or larger than the prefix length in the network statement.

```
config network
   edit <network number>
      set area <area>
      set prefix <network prefix> <mask>
```

For example:

```
(ospf) # config network
(network) # edit 1
(1) # set area 0.0.0.4
(1) # set prefix 10.1.1.0 255.255.255.0
```

5. Configure the OSPF interface

Configure interface-related OSPF settings. Enter a descriptive name for the OSPF interface name. Use the `set interface` command to apply this configuration to a FortiSwitch interface:

```
config ospf-interface
   edit <ospf interface name>
      set interface <interface name>
      set priority <>
```

For example:

```
(ospf) # config ospf-interface
(ospf-interface) # edit oi1
(oi1) # set interface vlan40-p4
(oi1) # set priority 255
```

NOTE: The following values must match for an adjacency to form:

- area type and number
- interface subnet and mask
- hello interval
- dead interval

6. Redistribute non-OSPF routes

Redistribute non-OSPF routes (directly connected or static routes) within OSPF:

```
config redistribute { <name> | connected | rip | static }
   set status enable
   set metric <integer>
   set metric-type {1 | 2}
end
```

For example:

```
(ospf) # config redistribute connected
(connected) # set status enable
```

Check the OSPF configuration

The `get router info ospf` command has options to display different aspects of the OSPF configuration and status. For example:

```
get router info ospf neighbors
get router info ospf database
```

Example configuration

The following example shows a very simple OSPF network with one area. FortiSwitch1 has one OSPF interface to FortiSwitch 2:

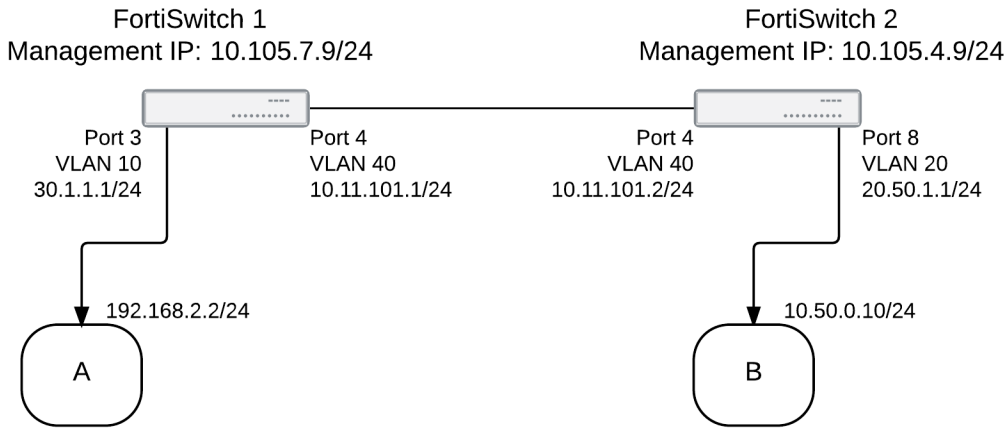

Configure system interfaces

These are the same configuration steps as for static routing.

Switch 1

```
config system interface
   edit vlan10-p3
      set ip 30.1.1.1 255.255.255.0
      set allowaccess ping https http ssh telnet
      set vlanid 10
   next
   edit vlan40-p4
      set ip 10.11.101.1 255.255.255.0
      set allowaccess ping https http ssh telnet
      set vlanid 40
end
config switch interface
   edit "port3"
      set native-vlan 10
   next
   edit "port4"
      set native-vlan 40
   next
end
```

Switch 2

```
config system interface
```

```
    edit vlan20-p8
       set ip 20.50.1.1 255.255.255.0
       set allowaccess ping https http ssh telnet
       set vlanid 20
    next
    edit vlan40-p4
       set ip 10.11.101.2 255.255.255.0
       set allowaccess ping https http ssh telnet
       set vlanid 40
end
config switch interface
    edit "port8"
       set native-vlan 20
    next
    edit "port4"
       set native-vlan 40
    next
end
```

Configure the OSPF router

Configure OSPF with the following:

1. Set the router ID.
2. Create the area.
3. Create the network (set network prefix and associate with an area).
4. Configure the OSPF interface.
5. Redistribute the routes.

Switch 1

```
config router ospf

    set router-id 10.11.101.1

    config area
       edit 0.0.0.0
       next
    end

    config network
       edit 1
          set area 0.0.0.0
          set prefix 10.11.101.0 255.255.255.0
       next
    end

    config ospf-interface
       edit "1"
          set cost 100
          set interface "vlan10"
          set priority 100
       next
    end
```

```
   config redistribute connected
      set status enable
   end

end
```

Switch 2

```
config router ospf
   set router-id 10.11.101.2

   config area
      edit 0.0.0.0
      next
   end

   config network
      edit 1
         set area 0.0.0.0
         set prefix 10.11.101.0 255.255.255.0
      next
   end

   config ospf-interface
      edit "1"
         set cost 100
         set interface "vlan10"
         set priority 100
      next
   end

   config redistribute connected
      set status enable
   end

end
```

Verify OSPF neighbors

```
get router info ospf neighbor all
```

Verify OSPF routes

```
get router info ospf route
```

RIP routing

NOTE: You must have an advanced features license to use RIP routing.

The Routing Information Protocol (RIP) is a distance-vector routing protocol that works best in small networks that have no more than 15 hops. Each router maintains a routing table by sending out its routing updates and by asking neighbors for their routes. RIP is relatively simple to configure on FortiSwitch units but slow to respond to network outages. RIP is better than static routing but less scalable than open shortest path first (OSPF).

FortiSwitch supports RIP version 1 and RIP version 2:

- RIP version 1 uses classful addressing and broadcasting to send out updates to router neighbors. It does not support different sized subnets or classless inter-domain routing (CIDR) addressing.
- RIP version 2 supports classless routing and subnets of various sizes. Router authentication supports MD5 and authentication keys. Version 2 uses multicasting to reduce network traffic.

RIP uses three timers:

- The update timer determines the interval between routing updates. The default setting is 30 seconds.
- The timeout timer is the maximum time that a route is considered reachable while no updates are received for the route. The default setting is 180 seconds. The timeout timer setting should be at least three times longer than the update timer setting.
- The garbage timer is the is the how long that the FortiSwitch advertises a route as being unreachable before deleting the route from the routing table. The default setting is 120 seconds.

You can enable bidirectional forwarding detection (BFD) with RIP. BFD is used to quickly locate hardware failures in the network. Routers running BFD communicate with each other, and, if a timer runs out on a connection, that router is declared to be down. BFD then communicates this information to RIP, and the routing information is updated.

For additional information about RIP routing, see the Routing Information Protocol (RIP) section of the FortiOS Handbook.

This chapter covers the following topics:

Terminology

Access list: A list of IP addresses and the action to take for each one. Access lists provide basic route and network filtering.

Active RIP interface: Each RIP router sends and receives updates by actively communicating with its neighbors.

Keychain: A list of one or more authentication keys including its lifetime, which is how long each key is valid.

Metric: RIP uses hop count as the metric for choosing the best route. A hop count of 1 represents a network that is connected directly to the FortiSwitch. A hop count of 16 represents a network that cannot be reached.

Passive RIP interface: The RIP router listens to updates from other routers but does not send out updates. A passive RIP interface reduces network traffic.

Prefix list: A more powerful prefix-based filtering mechanism. A prefix is an IP address and netmask.

Split horizon: A way to avoid routing loops.

Configuring RIP

NOTE: You must create a keychain first before you can use the MD5 authentication mode with RIP version 2.

To add a new keychain using the CLI:

```
config router key-chain
   edit <keychain identifier>
   next
end
```

Using the Web-based manager:

1. Create a switched virtual interface (SVI). See Configuring a switched virtual interface on page 143.
2. Go to **Router > Router > RIP**.
3. Select whether you want to use RIP version 1 or RIP version 2 and click **Apply**. RIP version 2 is the default.
4. If you have a complex configuration, select the appropriate options under **Advanced Options**.
5. Enter an IP address and netmask for your RIP network, separated with a slash, and click **Add**. For example, 172.168.200.0/255.255.255.0. **NOTE:** Select an IP address for a network that includes all SVIs that you want to use. You can configure multiple network ranges to cover all SVIs that will be using RIP routing.
6. To enable interface-specific features (such as authentication and the RIP version to send and receive routing updates), select the appropriate options under **Interfaces**.

Using the CLI:

```
config router rip
   config network
      edit <network identifier>
         set prefix <network prefix> <mask>
      next
   end
end
```

Checking the RIP configuration

The `get router info rip` command has options to display different aspects of the RIP configuration and status. For example, there are options to display the RIP general information and the RIP database:

```
get router info rip status
get router info rip database
```

Example configuration

The following example shows a very simple RIP network:

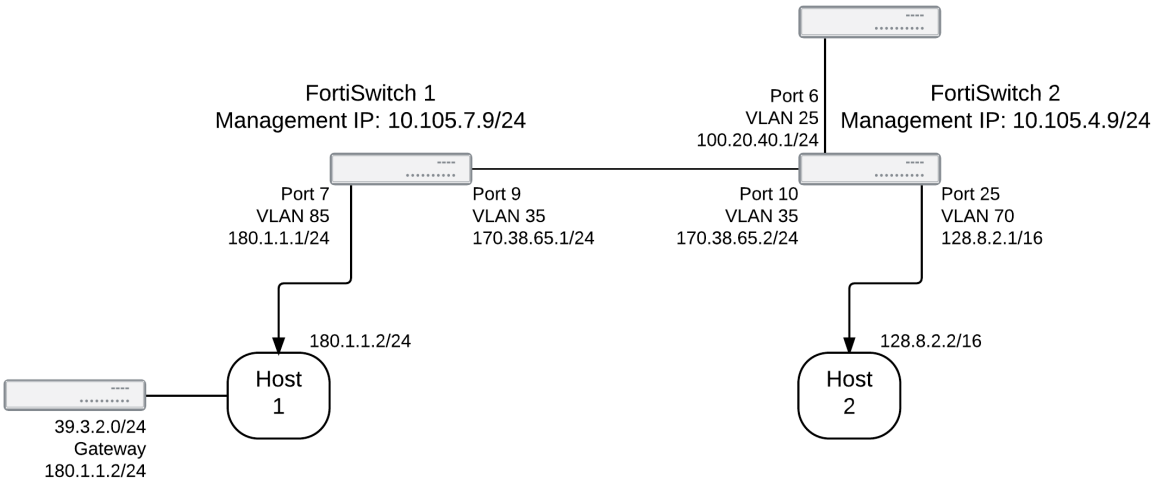

Switch 1: Configure the switch interface

```
config switch interface
   edit "port9"
      set allowed-vlans 35
   next
   edit "port7"
      set allowed-vlans 85
   next
end
```

Switch 1: Configure the system interface

```
config system interface
   edit "vlan35"
      set ip 170.38.65.1/24
      set allowaccess ping https http ssh snmp telnet
      set vlanid 35
   next
   edit "vlan85"
      set ip 180.1.1.1/24
      set allowaccess ping https http ssh snmp telnet
      set vlanid 85
   next
end
```

Switch 1: Configure the RIP router; add authentication between FortiSwitch 1 and FortiSwitch 2

```
config router rip
   config network
```

```
      edit 1
         set prefix 170.38.65.0/24
      next
      edit 2
         set prefix 180.1.1.0/24
      next
   end
   config interface
      edit "vlan35"
         set auth-mode text
         set auth-string simplepw1
      next
   end
end
```

Switch 1: Add a static route and redistribute it

```
config router static
   edit 1
      set dst 39.3.2.0 255.255.255.0
      set gateway 180.1.1.2
   next
end

config router rip
   config redistribute "static"
      set status enable
   next
end
```

Switch 2: Configure the switch interface

```
config switch interface
   edit "port10"
      set allowed-vlans 35
   next
   edit "port25"
      set allowed-vlans 70
   next
end
```

Switch 2: Configure the system interface

```
config system interface
   edit "vlan35"
      set ip 170.38.65.2/24
      set allowaccess ping https http ssh snmp telnet
      set vlanid 35
   next
   edit "vlan70"
      set ip 128.8.2.1/16
      set allowaccess ping https http ssh snmp telnet
      set vlanid 70
   next
end
```

Switch 2: Configure the RIP router; add authentication between FortiSwitch 1 and FortiSwitch 2

```
config router rip
   config network
      edit 1
         set prefix 170.38.65.0/24
      next
      edit 2
         set prefix 128.8.0.0/16
      next
   end
   config interface
      edit "vlan35"
         set auth-mode text
         set auth-string simplepw1
      next
   end
end
```

Switch 2: Add a connected route and redistribute it

```
config switch interface
   edit "port6"
      set allowed-vlans 25
   next
end
config system interface
   edit "vlan25"
      set ip 100.20.40.1/24
      set allowaccess ping https http ssh snmp telnet
      set vlanid 25
   next
end

config router rip
   config redistribute "connected"
      set status enable
   next
end
```

VRRP

NOTE: You must have an advanced features license to use VRRP.

The Virtual Router Redundancy Protocol (VRRP) uses virtual routers to control which physical routers are assigned to an access network. A VRRP group consists of a master router and one or more backup routers that share a virtual IP address. If the master router fails, the VRRP automatically assigns one of the backup routers without affecting network traffic. When the failed router is functioning again, it becomes the master router again. VRRP provides this redundancy without user intervention or additional configuration to any of the devices on the network.

To create a VRRP group, you need to create a VRRP virtual MAC address, which is a shared MAC address adopted by the VRRP master. The VRRP virtual MAC address feature is disabled by default. You must enable the VRRP virtual MAC address feature on all members of a VRRP group.

The VRRP master router sends VRRP advertisement messages to the backup routers. When the VRRP master router fails to send advertisement messages, the backup router with the highest priority takes over as the master router.

For additional information about VRRP, see the VRRP section of the FortiOS Handbook.

This chapter covers the following topics:

Configuring VRRP

Using the Web-based manager:

1. Go to **System > Network > Interface**.
2. Select an interface and click **Edit**.
3. (Optional) Select the **VRRP Virtual MAC** check box.
4. Click **Add** to add a virtual router.
5. Enter the unique virtual router identifier.
6. Enter the VRRP group number.
7. Enter the priority. If the highest priority value of 255 is entered, the virtual router becomes the master router.
8. Select the **Preempt** check box if you want the router to preempt the master virtual router if the priority changes.
9. Enter the virtual IP address that will be shared across the VRRP group.
10. Enter one or two IP addresses that the master router must track. The maximum number of IP addresses is two. If these IP addresses cannot be reached by the master router, the priority of the master router changes to 0.
11. Click **OK**.
12. Click **Add** to add each additional virtual router.
13. After filling in the fields for each additional virtual router, click **OK**.

Using the CLI:

```
config system interface
   edit <VLAN name>
      set ip <IP address> <netmask>
      set allowaccess <access_types>
      set vrrp-virtual-mac enable
         config vrrp
            edit <VRRP router identifier>
               set priority <priority number>
               set vrgrp <VRRP group number>
               set vrip <virtual IP address>
            next
         end
      set snmp-index <index number>
      set vlanid <VLAN identifier>
      set interface "internal"
   next
end
```

Example of configuring VRRP:

```
config system interface
   edit "vlan-8"
      set ip 10.10.10.1 255.255.255.0
      set allowaccess ping https http ssh
      set vrrp-virtual-mac enable
         config vrrp
            edit 5
               set priority 255
               set vrgrp 50
               set vrip 11.1.1.100
            next
            edit 6
               set priority 200
               set vrgrp 50
               set vrip 11.1.1.100
            next
            edit 7
               set priority 150
               set vrgrp 50
               set vrip 11.1.1.100
            next
         end
      set snmp-index 20
      set vlanid 8
      set interface "internal"
   next
end
```

Checking the VRRP configuration

Use the `get router info vrrp` command to display the VRRP status:

```
get router info vrrp
```

Users and user groups

FortiSwitch provides authentication mechanisms to control user access to the system (based on the user group associated with the user). The members of user groups are user accounts. Local users and peer users are defined on the FortiSwitch. User accounts can also be defined on remote authentication servers.

This section describes how to configure local users and peer users and how to configure user groups. For information about configuring the authentication servers, see Remote authentication servers on page 33.

This chapter covers the following topics:

Users

A user account consists of a user name, password, and potentially other information, configured in a local user database or on an external authentication server.

Users can access resources that require authentication only if they are members of an allowed user group.

Local and remote users are defined in **System** > **User** > **User Definition**.

```
config user local
   edit <user_name>
      set ldap-server <server_name>
      set passwd <password_string>
      set radius-server <server_name>
      set tacacs+-server <server_name>
      set status {enable | disable}
      set type <auth-type>
   end
```

Field	Description
user_name	Identifies the user
password_string	A password for the local user
ldap-server <server_name>	To authenticate this user using a password stored on a remote authentication server, select the type of server and then select the server from the list. You can select only a server that has already been added to the FortiSwitch configuration.
radius-server <server_name>	
tacacs+-server <server_name>	
status	Enable or disable this user.

User groups

A user group contains a list of local and remote users.

Security policies allow access to specified user groups only. This restricted access enforces Role Based Access Control (RBAC) to your organization's network and its resources. Users must be in a group and that group must be part of the security policy.

```
config user group
   edit <groupname>
      set authtimeout <timeout>
      set group-type <grp_type>
      set http-digest-realm <attribute>
      set member <names>
      config match
         edit <match_id>
            set group-name <gname_str>
            set server-name <srvname_str>
      end
   end
```

The following table describes the parameters:

Field	Description
groupname	Identifies the user group.
authtimeout <timeout>	Sets the authentication timeout for the user group. The range is 1 to 480 minutes. If this field is set to 0, the global authentication timeout value is used.
group-type <grp_type>	Enter the group type. <grp_type> determines the type of users and is one of the following: • **firewall** - FortiSwitch users defined in user local, user ldap, or user radius • **fsso-service** - Directory Service users
http-digest-realm <attribute>	Enter the realm attribute for MD5-digest authentication.
member <names>	Enter the names of users, peers, LDAP servers, or RADIUS servers to add to the user group. Separate the names with spaces. To add or remove names from the group, you must re-enter the whole list with the additions or deletions required.
config match fields	
<match_id>	Enter an ID for the entry.
group-name <gname_str>	Identifies the matching group on the remote authentication server.
server-name <srvname_str>	Specifies the remote authentication server.

802.1x authentication

To control network access, FortiSwitch supports IEEE 802.1x authentication. A supplicant connected to a port on the switch must be authenticated by a RADIUS/Diameter server to gain access to the network. The supplicant and the authentication server communicate using the switch using the EAP protocol. FortiSwitch supports EAP-PEAP, EAP-TTLS, EAP-TLS, and EAP-MD5.

To use the RADIUS server for authentication, you must configure the server before configuring the users or user groups on the FortiSwitch.

FortiSwitch implements MAC-based authentication. The switch saves the MAC address of each supplicant's device. The switch provides network access only to devices that have successfully been authenticated.

You can enable the MAC Authentication Bypass (MAB) option for devices (such as network printers) that cannot respond to the 802.1x authentication request. With MAB enabled on the port, the system will use the device MAC address as the user name and password for authentication.

Optionally, you can configure a guest VLAN for unauthorized users. Alternatively, you can specify a VLAN for users whose authentication was unsuccessful.

NOTE: You cannot configure persistent static MAC addresses with 802.1x authentication.

This chapter covers the following topics:

Dynamic VLAN assignment

You can configure the RADIUS server to return a VLAN in the authentication reply message.

1. On the FortiSwitch, select port-based authentication or MAC-based authentication and a security group.
2. On the RADIUS server, configure the attributes.

Using the Web-based manager:

1. Go to **Switch > Interface > Interface**.

2. Select a port and click **Edit**.

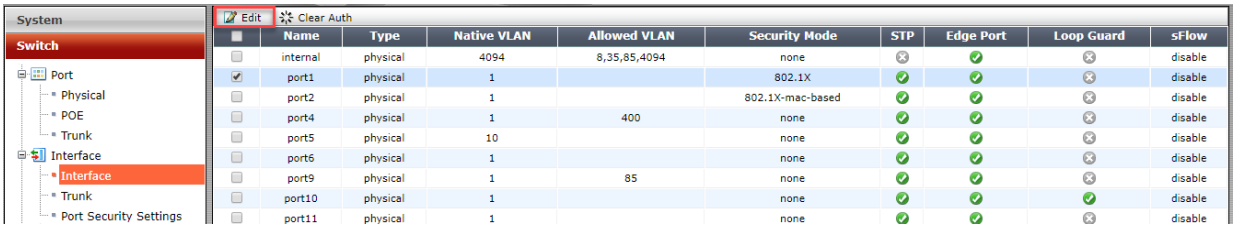

3. Select **802.1X** for port-based authentication or select **802.1X-MAC-based** for MAC-based authentication.

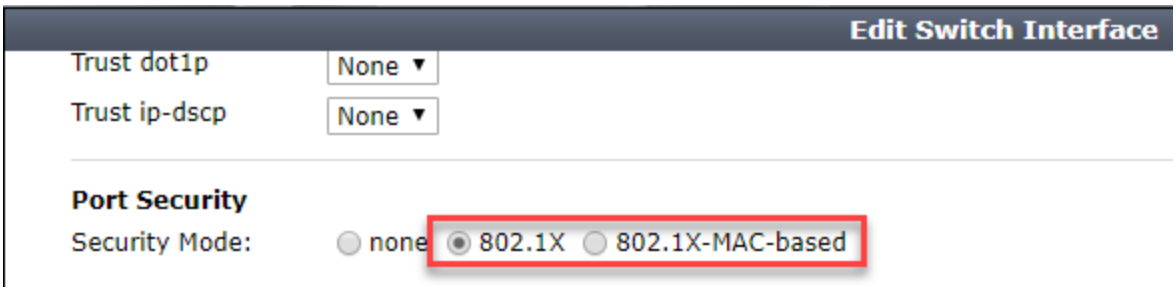

4. Select a security group.

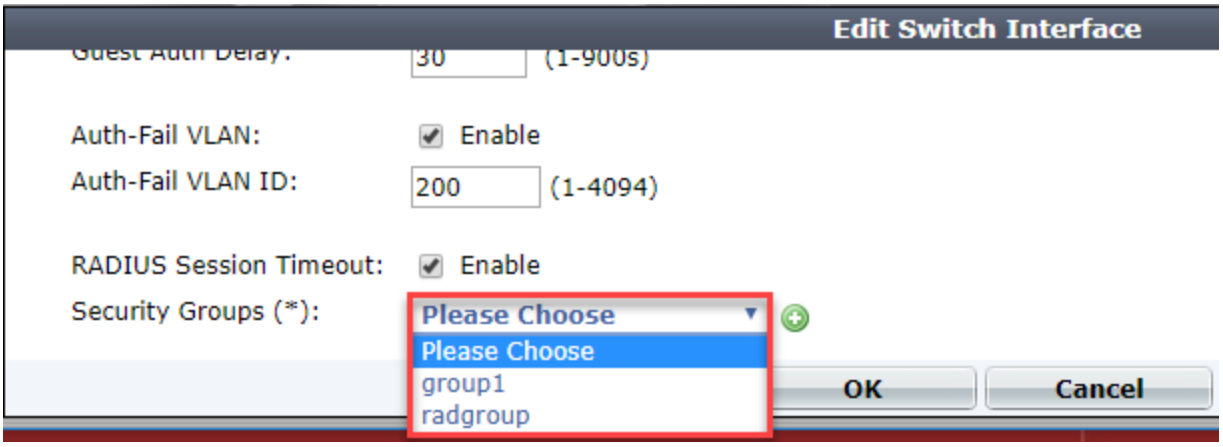

5. Click **OK**.

Using the CLI:

To select port-based authentication and the security group on the FortiSwitch:

```
config switch interface
   edit <interface_name>
      config port-security
         set port-security-mode 802.1X
      end
      set security-groups <security-group-name>
   end
```

The FortiSwitch will change the native VLAN of the port to that of the VLAN from the server.

To select MAC-based authentication and the security group on the FortiSwitch:

```
config switch interface
   edit <interface_name>
      config port-security
         set port-security-mode 802.1X-mac-based
      end
      set security-groups <security-group-name>
   end
```

Here, the switch assigns the returned VLAN only to this user's MAC address. The native VLAN of the port remains unchanged.

Use the following configuration command to view the MAC-based VLAN assignments:

```
diagnose switch vlan assignment mac list [sorted-by-mac | sorted-by-vlan]
```

Configure the following attributes in the RADIUS server:

- Tunnel-Private-Group-Id—10 (vlanid)
- Tunnel-Medium-Type—IEEE-802(6)
- Tunnel-Type—VLAN (13)

MAC authentication bypass (MAB)

Devices such as network printers, cameras, and sensors might not support 802.1x authentication. If you enable the MAB option on the port, the system will use the device MAC address as the user name and password for authentication.

MAB retries authentication three times before the device is assigned to a guest VLAN for unauthorized users. By default, reauthentication is disabled. Use the following commands if you want to change the default behavior:

```
config switch global
   config port-security
      set mab-reauth enable
   end
```

You must provision the RADIUS server to authenticate the devices that use MAB, either by adding the MAC addresses as regular users or by implementing additional logic to resolve the MAC addresses in a network inventory database.

The following flowchart shows the FortiSwitch 802.1x port-based authentication with MAB enabled:

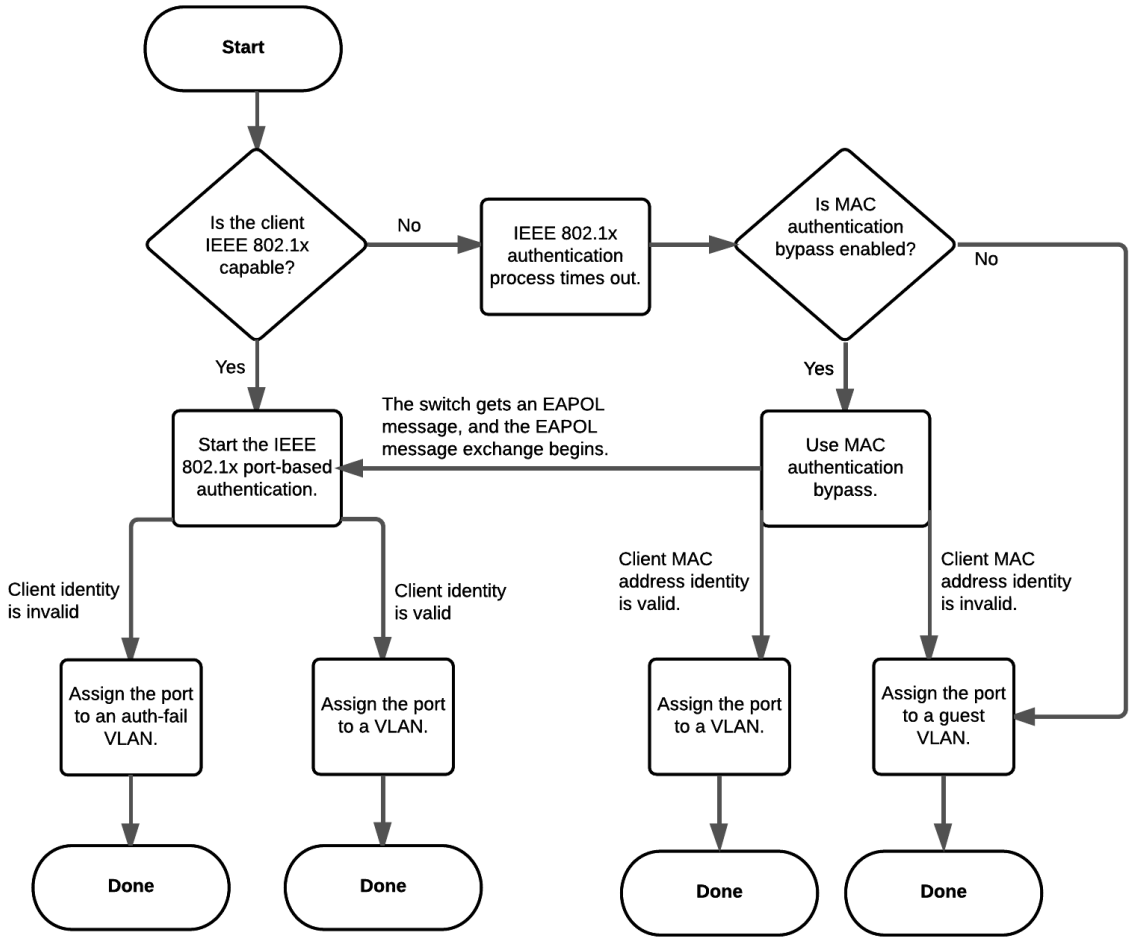

The following flowchart shows the FortiSwitch 802.1x MAC-based authentication with MAB enabled:

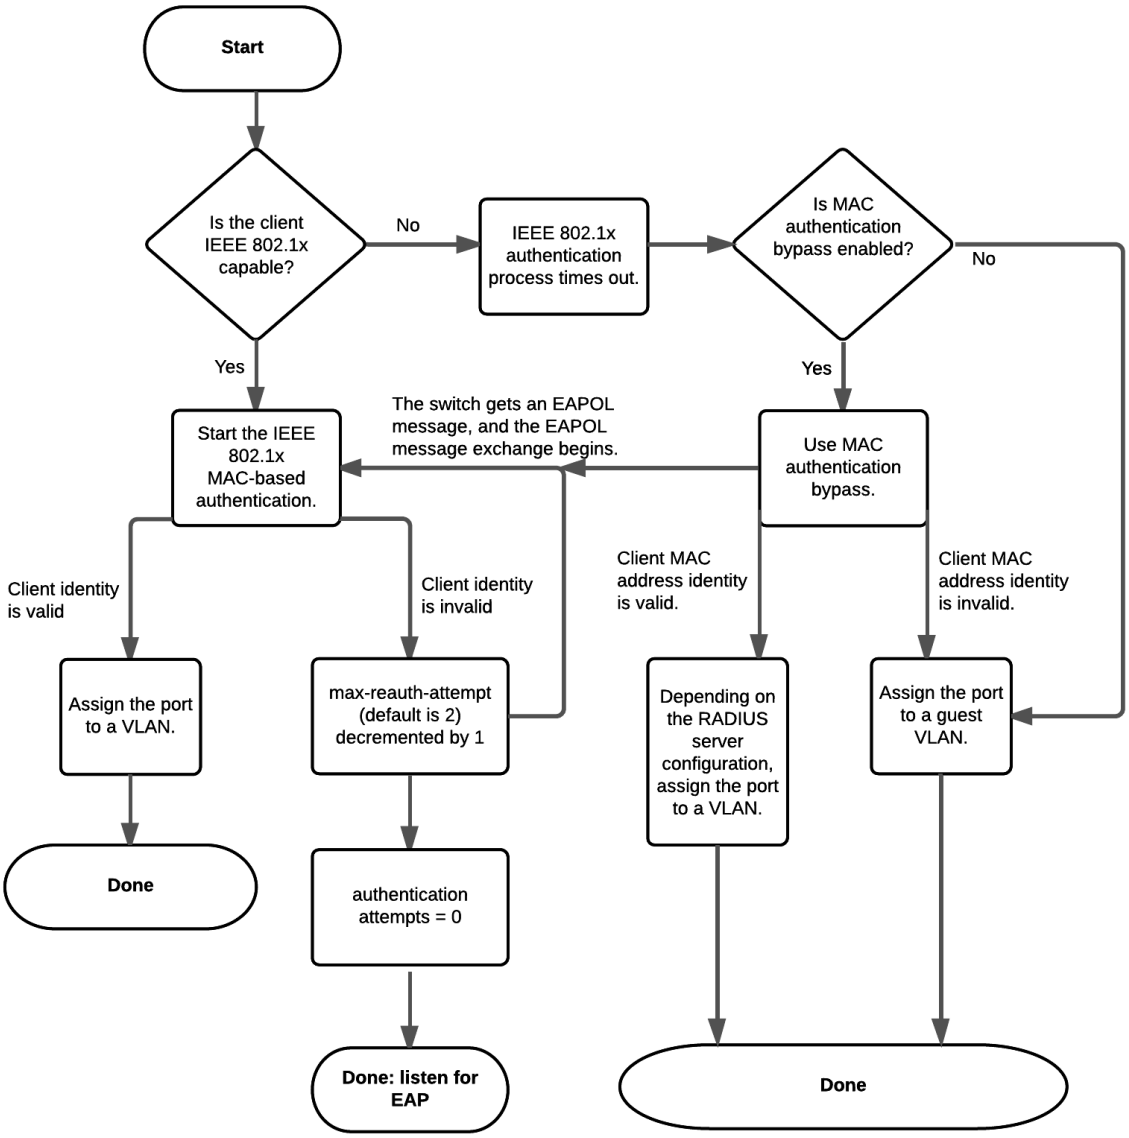

Configuring global settings

Using the Web-based manager:

1. Go to **Switch > Interface > Port Security Settings**.

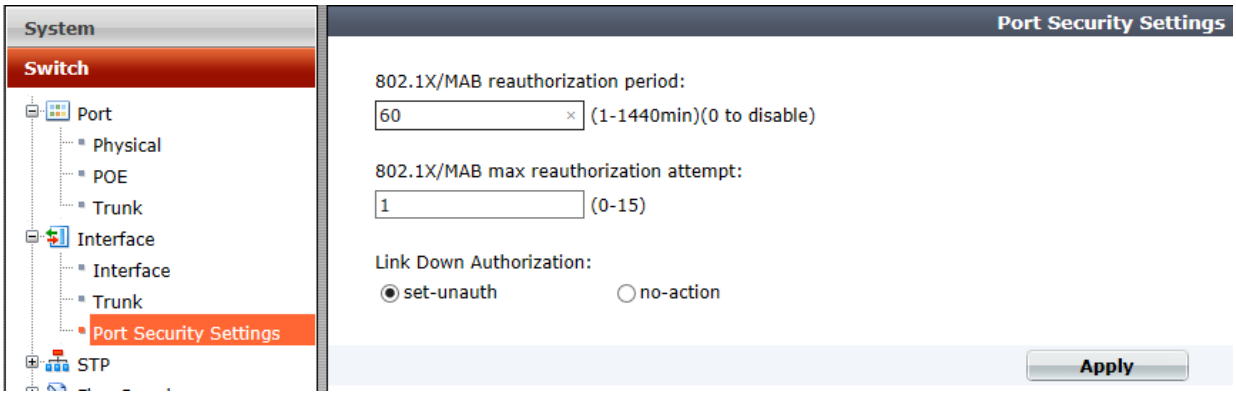

2. In the 802.1X/MAB reauthorization period field, enter the number of minutes before the system requires the device to reauthenticate.
3. In the 802.1X/MAB max authorization attempt field, enter the maximum number of times that the system will try to reauthorize the session.
4. Select **set-unauth** to revert all devices to the unauthenticated state if the link goes down or select **no-action** if reauthentication is unnecessary if the link goes down.
5. Click **Apply**.

Using the CLI:

```
config switch global
   config port-security
      set reauth-period <0-1440>
      set max-reauth-attempt <0-15>
      set link-down-auth {no-action | set-unauth}
```

NOTE: Changes to global settings only take effect when new 802.1x/MAB sessions are created.

Variable	Description
reauth-period	This setting defines how often the device needs to reauthenticate (that is, if a session remains active beyond this number of minutes, the system requires the device to reauthenticate). The default value is 60 minutes. Set the value to 0 to disable reauthentication.
max-reauth-attempt	If 802.1x authentication fails, this setting caps the number of reattempts that the system will initiate. The range is from 0 to 15 where "0" translates to forever (fail causes a log message). The default value is 3.
link-down-auth	If a link goes down, this setting determines whether the impacted devices must reauthenticate. Set the value to `no-action` if reauthenication is unnecessary. Set the value to `set-unauth` to revert all devices to the unauthenticated state. Each device must reauthenticate. The default is `set-unauth`.

Configuring the 802.1x settings on an interface

Using the Web-based manager:

1. Go to **Switch > Interface >Interface**.
2. Select a port and click **Edit**.

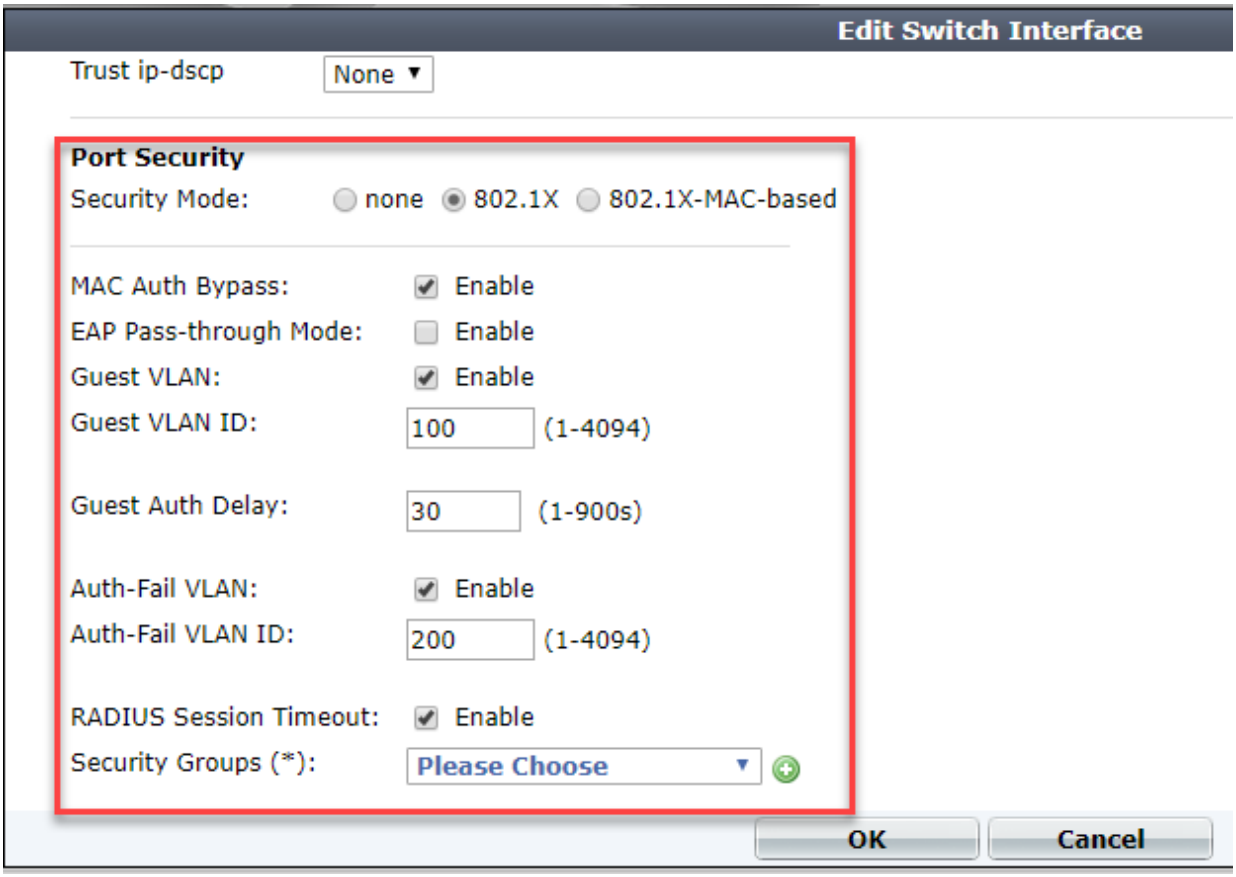

3. Select **802.1X** for port-based authentication or select **802.1X-MAC-based** for MAC-based authentication.
4. Select **MAC Auth Bypass Enable**.
5. Select **Guest VLAN Enable** if you want to assign a VLAN to unauthorized users. If you select **Enable**, enter the VLAN identifier in the Guest VLAN ID field.
6. In the Guest Auth Delay field, enter the number of seconds for an unauthorized user to have access as a guest before authorization fails.
7. Select **Auth-Fail VLAN Enable** if you want to assign a VLAN to users who attempted to authenticate but failed to provide valid credentials. If you select **Enable**, enter the VLAN identifier in the Auth-Fail VLAN ID field.
8. If you want to use the RADIUS-provided reauthentication time, select **RADUS Session Timeout Enable**.
9. If you selected the 802.1X or 802.1X-MAC-based security mode, you must select a security group.
10. Click **OK**.

Using the CLI:

```
config switch interface
   edit <port>
      config port-security
         set port-security-mode {none | 802.1X | 8021X-mac-based}
         set mac-auth-bypass {enable | disable}
         set guest-vlan {enable | disable}
         set guest-vlanid <vlanid>
         set guest-auth-delay <integer>
         set auth-fail-vlan {enable | disable}
         set auth-fail-vlanid <vlanid>
         set radius-timeout-overwrite {enable | disable}
      end
      set security-groups <security-group-name>
   end
```

Variable	Description
port-security-mode	Set the security mode for the port. None (no security) is the default. Set the security mode to `802.1X` for port-based authentication or `802.1X-mac-based` for MAC-based authentication. If you change the security mode from none, you must set the security group with the `set security-groups` command.
mac-auth-bypass	Enable the feature. The default is disable.
guest-vlan and auth-fail-vlan	The system assigns the guest-vlan to unauthorized users. After the system assigns the auth-fail-vlan to users who attempted to authenticate but failed to provide valid credentials. If you enable either `guest-vlan` or `auth-fail-vlan`, you must configure the corresponding VLAN ID (otherwise, the configuration save attempt will fail when you enter `next` or `end`).
guest-auth-delay	Time when an authorization fails after the guest is applied. In seconds ranging from 60 to 900. The default is 120.

Variable	Description
radius-timeout-overwrite	This setting specifies whether to use the RADIUS-provided re-authentication timeout. If the setting is enabled, the port uses the local timeout (see Configuring global settings on page 174). If the setting is disabled, the system uses the value of the RADIUS Access-Accept message Session-Timeout attribute to determine the duration of the session. It uses the Termination-Action value to determine the device action when the session's timer expires. If the Termination-Action attribute is present and its value is RADIUS-Request, the device port re-authenticates the host. If the Termination-Action attribute is not present, or its value is Default, the device port terminates the session. If the device port is configured to use the RADIUS-supplied timeout, but the Access-Accept message does not include a Session-Timeout attribute, the device port never re-authenticates the supplicant.
security-groups <security-group-name>	Enter the security group name if you are using port-based authentication or MAC-based authentication.

Viewing the 802.1x details

Using the Web-based manager:

1. Go to **System >Switch >Monitor >802.1x Status**.

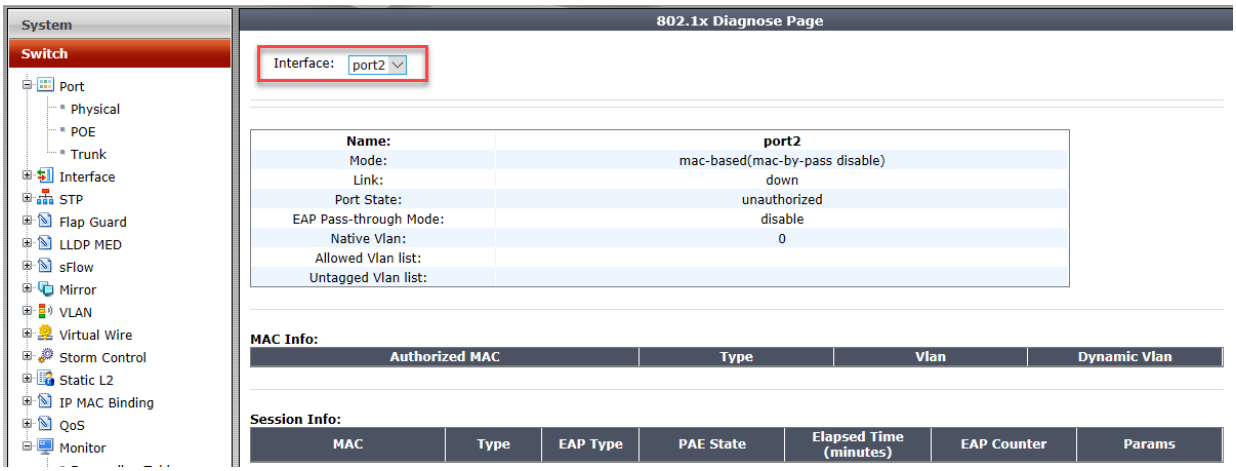

2. Select the appropriate interface.

Using the CLI:

Use the following command to show diagnostics on one or all ports:

```
diagnose switch 802-1x status [<port>]
```

```
port3 : Mode: port-based (MAC by-pass disable)
       Link: Link up
       Port State: authorized
       Dynamic Authorized Vlan: 10
       Native vlan: 10
       Allowed vlan list: 1-10
       Untagged vlan list:
       Guest vlan:
       AuthFail vlan:

       Sessions info:
       STA=00:24:9b:1b:20:65 Type=802.1X EAP PEAP state=AUTHENTICATED

port7 : Mode: mac-based (mac-by-pass disable)
       Link: Link up
       Port State: authorized ( )
       EAP pass-through mode : Enable
       Native Vlan : 1
       Allowed Vlan list: 1
       Untagged Vlan list: 1
       Guest VLAN :

       Client MAC Type Vlan Dynamic-Vlan
       0a:0a:0b:0b:0a:0a 802.1x 1 0
       0a:0a:0b:0b:0a:09 802.1x 1 0
       0a:0a:0b:0b:0a:08 802.1x 1 0
       0a:0a:0b:0b:0a:07 802.1x 1 0
       0a:0a:0b:0b:0a:06 802.1x 1 0
       0a:0a:0b:0b:0a:05 802.1x 1 0
       0a:0a:0b:0b:0a:04 802.1x 1 0
       0a:0a:0b:0b:0a:03 802.1x 1 0
       0a:0a:0b:0b:0a:02 802.1x 1 0
       0a:0a:0b:0b:0a:01 802.1x 1 0

       Sessions info:
       0a:0a:0b:0b:0a:0a Type=802.1x,MD5,state=AUTHENTICATED,etime=2,eap_cnt=3 param-
s:reAuth=3600
       0a:0a:0b:0b:0a:09 Type=802.1x,MD5,state=AUTHENTICATED,etime=2,eap_cnt=3
params:reAuth=3600
       0a:0a:0b:0b:0a:08 Type=802.1x,MD5,state=AUTHENTICATED,etime=2,eap_cnt=3 param-
s:reAuth=3600
       0a:0a:0b:0b:0a:07 Type=802.1x,MD5,state=AUTHENTICATED,etime=2,eap_cnt=3
params:reAuth=2896
       0a:0a:0b:0b:0a:06 Type=802.1x,MD5,state=AUTHENTICATED,etime=2,eap_cnt=3 param-
s:reAuth=3600
       0a:0a:0b:0b:0a:05 Type=802.1x,MD5,state=AUTHENTICATED,etime=2,eap_cnt=3
params:reAuth=3600
       0a:0a:0b:0b:0a:04 Type=802.1x,MD5,state=AUTHENTICATED,etime=2,eap_cnt=3 param-
s:reAuth=3600
       0a:0a:0b:0b:0a:03 Type=802.1x,MD5,state=AUTHENTICATED,etime=2,eap_cnt=3
params:reAuth=3600
       0a:0a:0b:0b:0a:02 Type=802.1x,MD5,state=AUTHENTICATED,etime=2,eap_cnt=3
```

```
params:reAuth=3600
       0a:0a:0b:0b:0a:01 Type=802.1x,MD5,state=AUTHENTICATED,etime=2,eap_cnt=3 param-
s:reAuth=3600h=120
```

Clearing port authorizations

Use the following command to manually flush all authorizations on a given port:

```
execute 802-1x clear interface <port>
```

Authenticating users with a RADIUS server

Using the Web-based manager:

1. Define the RADIUS server:
   a. Go to **System > Authentication > RADIUS Servers**.
   b. Click **Create New**.

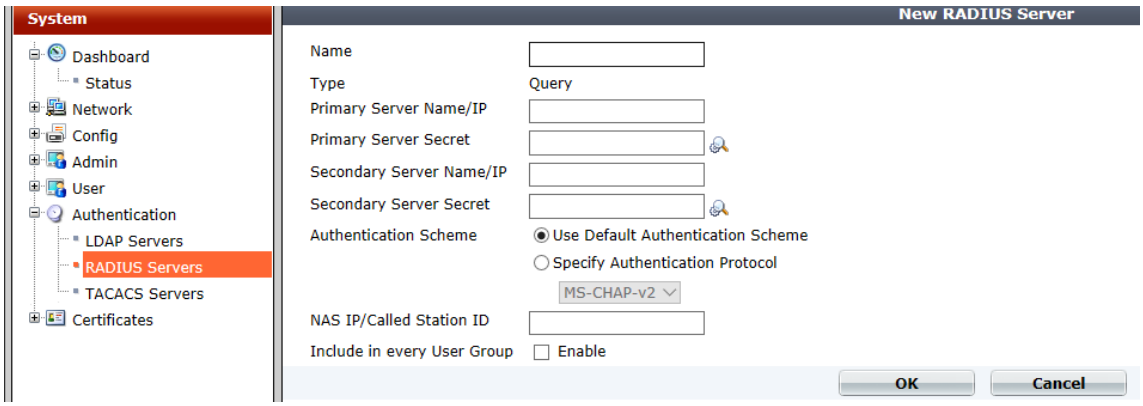

   c. In the Name field, enter a name for the RADIUS server.
   d. In the Primary Server Name/IP field, enter the IP address for the RADIUS server.
   e. In the Primary Server Secret field, enter a password to use as a RADIUS key.
   f. Click **OK**.
2. Create a user group:
   a. Go to **System > User > User Groups**.

b. Click **Create New**.

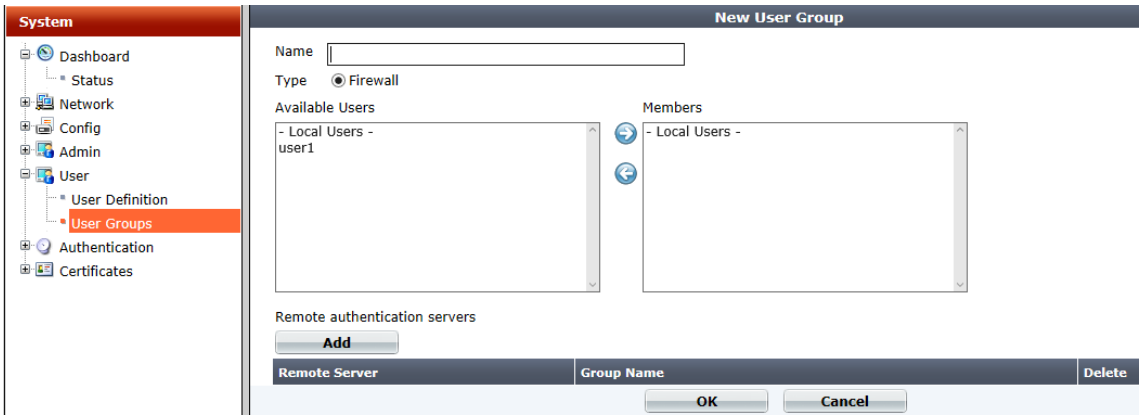

c. In the new field, enter a name for the user group.

d. Click **Add**.

e. Select the name of the RADIUS server that you configured in step 1.

f. Click **OK**.

3. Configure the port security:

a. Go to **Switch > Interface > Interface**.

b. Select a port and click **Edit**.

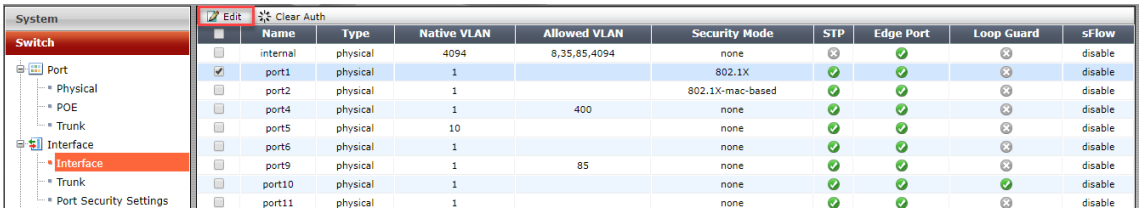

c. Select **802.1X** for port-based authentication or select **802.1X-MAC-based** for MAC-based authentication.

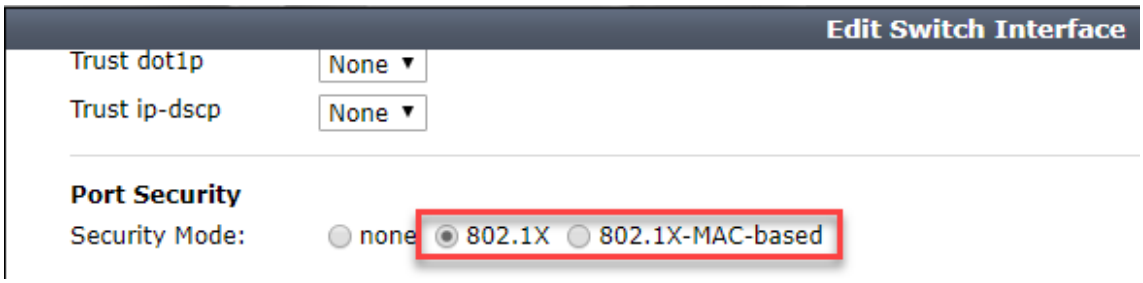

d. Select the user group that you configured in step 2.

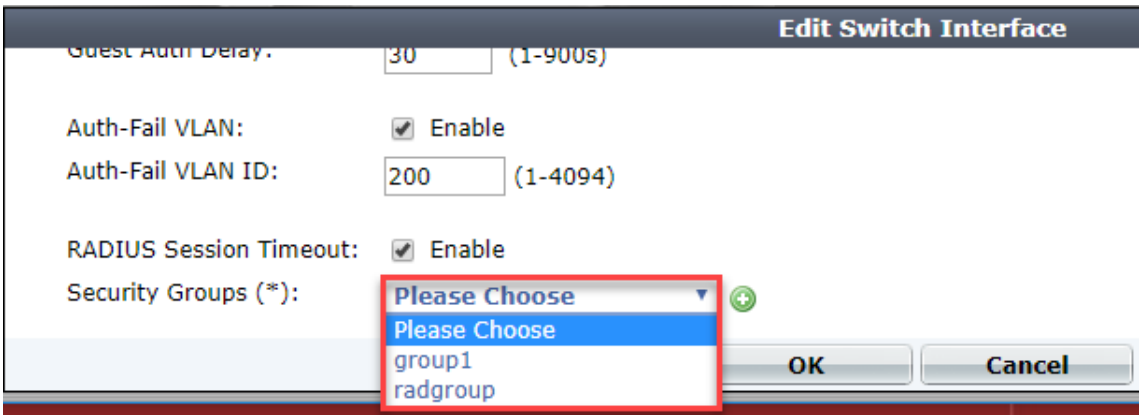

e. Click **OK**.

Using the CLI:

1. Define the RADIUS server:

```
config user radius
   edit <name>
      set server <address>
   end
end
```

2. Create a user group:

```
config user group
   edit <name>
      set member <list>
      config match
         edit 1
            set group-name <name>
            set server-name <name>
         end
      end
   end
end
```

3. Configure the switch interface for port-based 802.1x:

```
config switch interface
   edit <interface>
      config port-security
         set port-security-mode 802.1X
      end
      set security-groups <security-group-name>
   end
end
```

4. Configure the switch interface for MAC-based 802.1x:

```
config switch interface
   edit <interface>
      config port-security
         set port-security-mode 802.1X-mac-based
```

```
        end
        set security-groups <security-group-name>
    end
end
```

Example: RADIUS user group

Using the Web-based manager:

1. Define the RADIUS server:
   a. Go to **System > Authentication > RADIUS Servers**.
   b. Click **Create New**.
   c. In the Name field, enter **FortiAuthenticator**.
   d. In the Primary Server Name/IP field, enter **10.160.36.190**.
   e. In the Primary Server Secret field, enter **6rF7O4/Zf3p2TutNyeSjPbQc73QrS21wNDmNXd/rg9k6nTR6yMhBRsJGpArhIe6UOCb7b8InM3nrCeuVETr/a02LpILmIltBq5sUMCNqbR6zp2fS3r35Eyd3IIrzmve4Vusi52c1MrCqVhzzy2EfxkBrx5FhcRQWxStvnVt4+dzLYbHZ**.

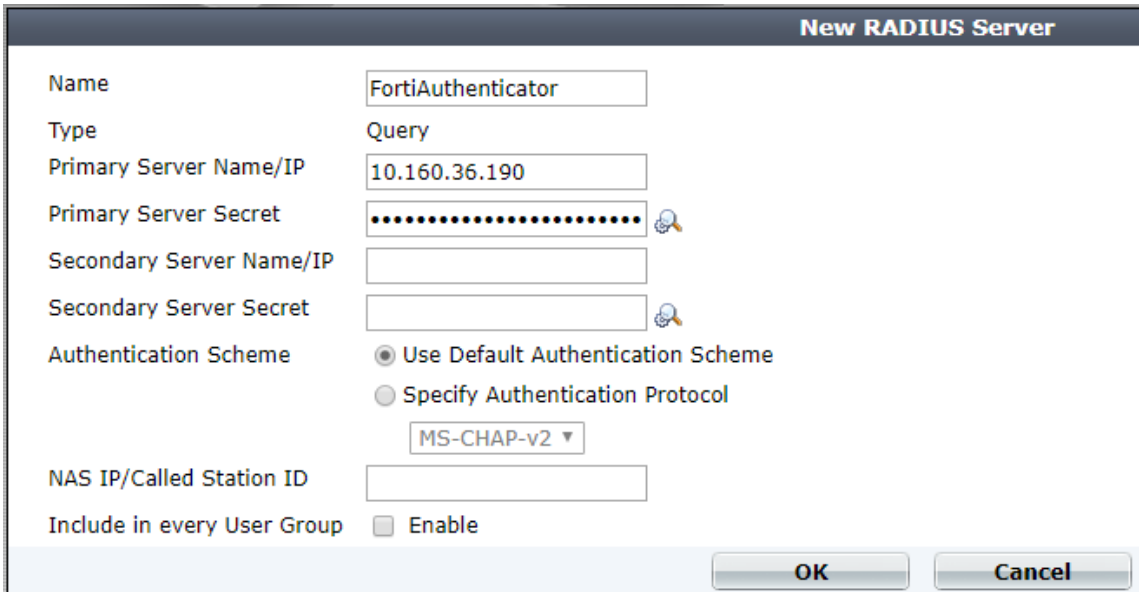

   f. Click **OK**.
2. Create a user group:
   a. Go to **System > User > User Groups**.
   b. Click **Create New**.
   c. In the Name field, enter **Radius_group**.
   d. Click **Add**.

e. Select **FortiAuthenticator** as the remote server.

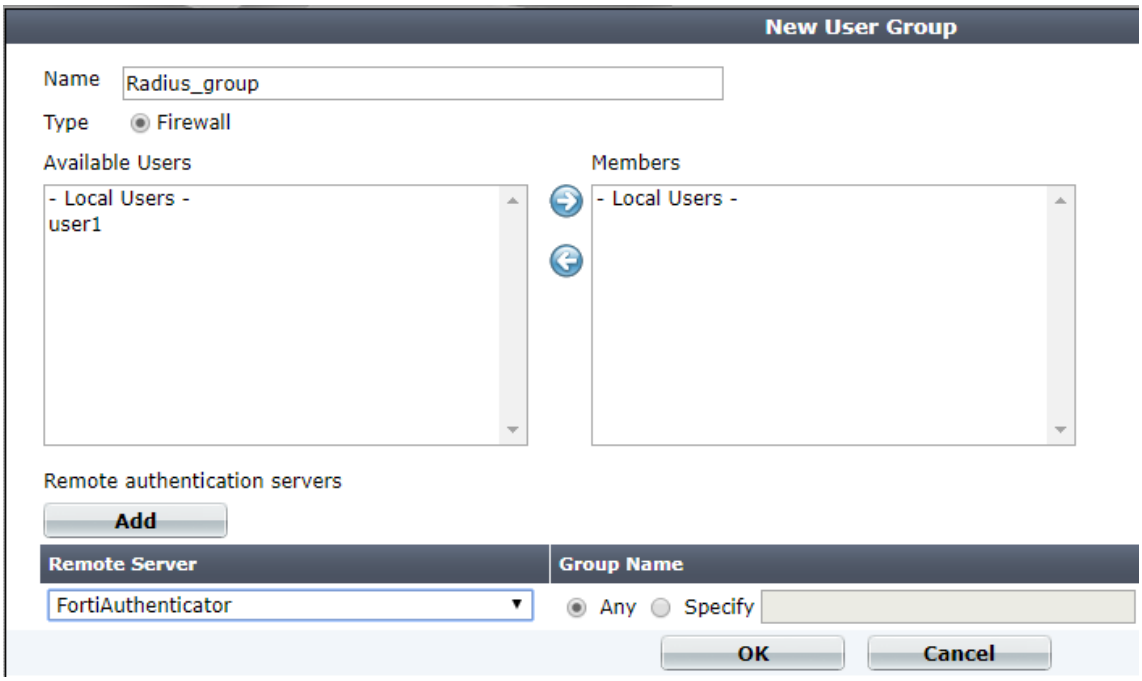

f. Click **OK**.

3. Configure the port security:
   a. Go to **Switch > Interface > Interface**.
   b. Select **port1** and click **Edit**.

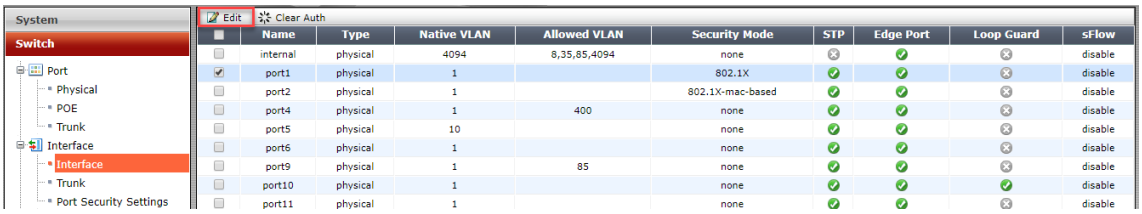

   c. In the Allowed VLAN field, enter **1**.
   d. Click **802.1X**.

e. Select **Radius_group**.

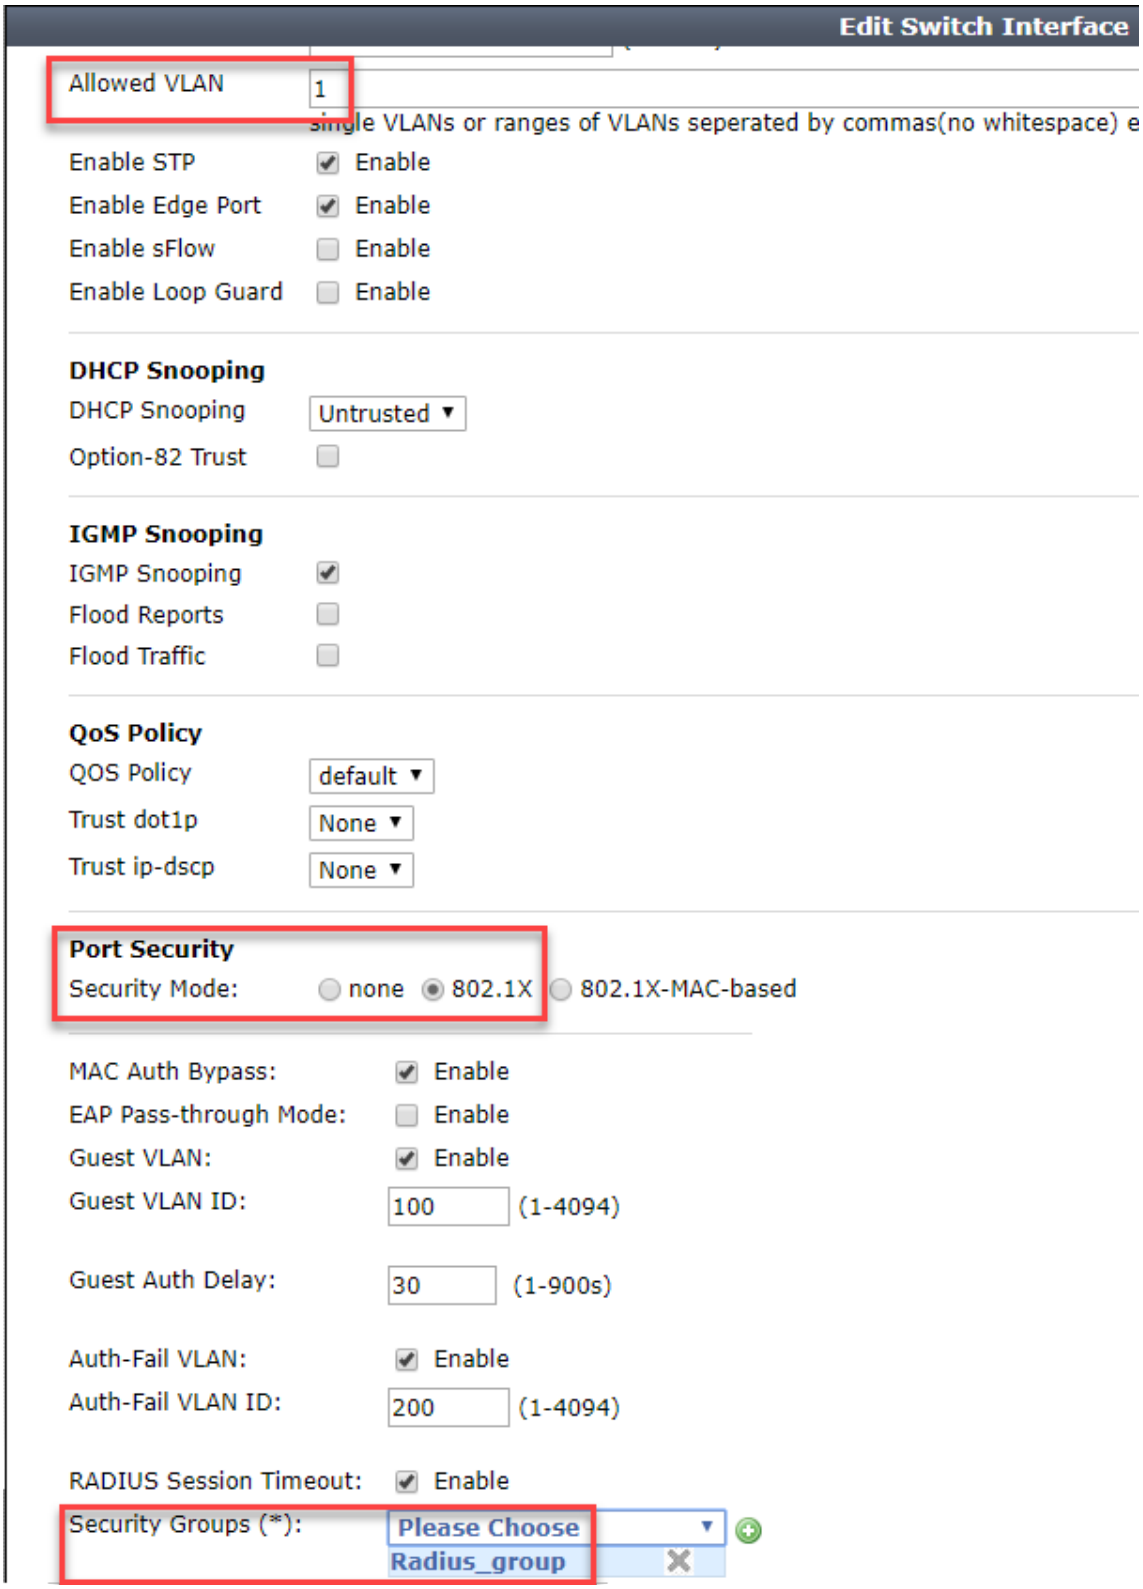

f. Click **OK**.

Using the CLI:

1. Define the RADIUS server:

```
config user radius
   edit "FortiAuthenticator"
      set secret ENC
          6rF7O4/Zf3p2TutNyeSjPbQc73QrS21wNDmNXd/rg9k6nTR6yMhBRsJGpArhle6UOCb7b8InM3n
          rCeuVETr/a02LpILmIltBq5sUMCNqbR6zp2fS3r35Eyd3IIrzmve4Vusi52c1MrCqVhzzy2Efxk
          Brx5FhcRQWxStvnVt4+dzLYbHZ
      set server "10.160.36.190"
   next
end
```

2. Create a user group:

```
config user group
   edit "Radius_group"
      set member "FortiAuthenticator"
   end
end
```

3. Configure the port security:

```
config switch interface
   edit "port1"
      set allowed-vlans 1
      set snmp-index 1
      config port-security
         set port-security-mode 802.1X
      end
      set security-groups "Radius_group"
   end
end
```

Example: dynamic VLAN

To assign VLAN dynamically for a port on which a user is authenticated, configure the RADIUS server attributes to return the VLAN ID when the user is authenticated. Assuming that the port security mode is set to 802.1X, the FortiSwitch will change the native VLAN of the port to the value returned by the server.

Ensure that the following attributes are configured on the RADIUS server:

- Tunnel-Private-Group-Id <integer> (the VLAN ID)
- Tunnel-Medium-Type IEEE-802 (6)
- Tunnel-Type VLAN (13)

Authenticating an admin user with RADIUS

If you want to use a RADIUS server to authenticate administrators, you must configure the authentication before you create the administrator accounts. Do the following:

1. Configure the FortiSwitch to access the RADIUS server.
2. Configure an administrator to authenticate with a RADIUS server and match the user secret to the RADIUS server

entry.

3. Create the RADIUS user group.

Using the Web-based manager:

1. Create a RADIUS system admin group:
   a. Go to **System > Admin > Administrators**.
   b. Click **Create New**.
   c. In the Administrator field, enter **RADIUS_Admins**.
   d. Click **Remote**.
   e. For the user group, select **Radius_group**.
   f. Select **Wildcard**.
   g. For the admin profile, select **super_admin**.

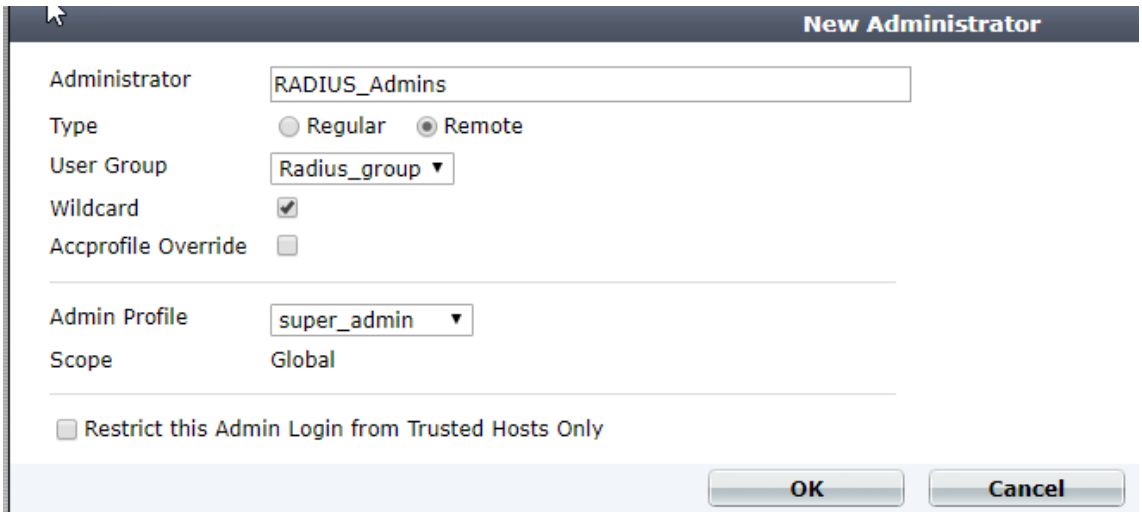

   h. Click **OK**.
2. Create a user:
   a. Go to **System > User > User Definition**.
   b. Click **Create New**.
   c. In the Name field, enter **RADIUS1**.
   d. Click **Password**.
   e. In the Password field, enter **6rF7O4/Zf3p2TutNyeSjPbQc73QrS21wNDmNXd/rg9k6nTR6yMhBRsJGpArhIe6UOCb7b8InM3nrCeuVETr/a02LpILmIltBq5sUMCNqbR6zp2fS3r35Eyd3IIrzmve4Vusi52c1MrCqVhzzy2EfxkBrx5FhcRQWxStvnVt4+dzLYbHZ**.

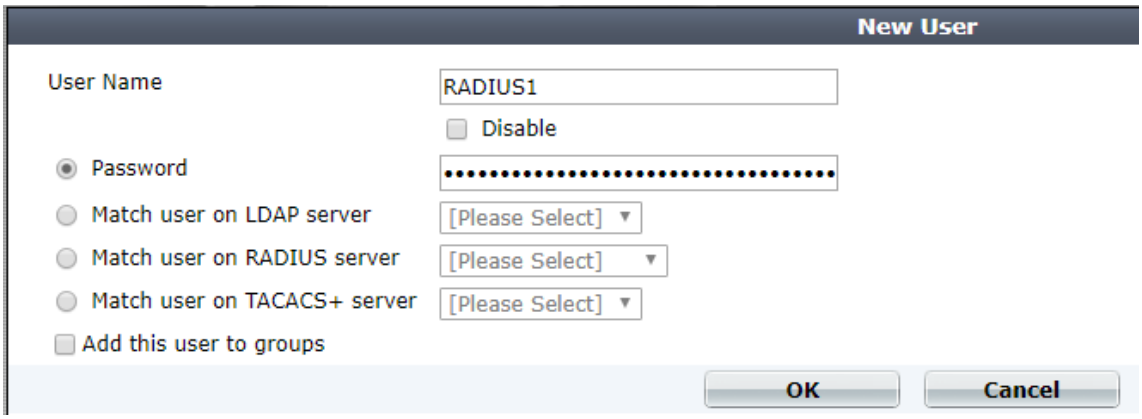

f. Click **OK**.

3. Create a user group:
   a. Go to **System > User > User Groups**.
   b. Click **Create New**.
   c. In the Name field, enter **RADIUS_Admins**.
   d. Select **RADIUS1** in the Available Users box and click the right arrow to move it to the Members box.

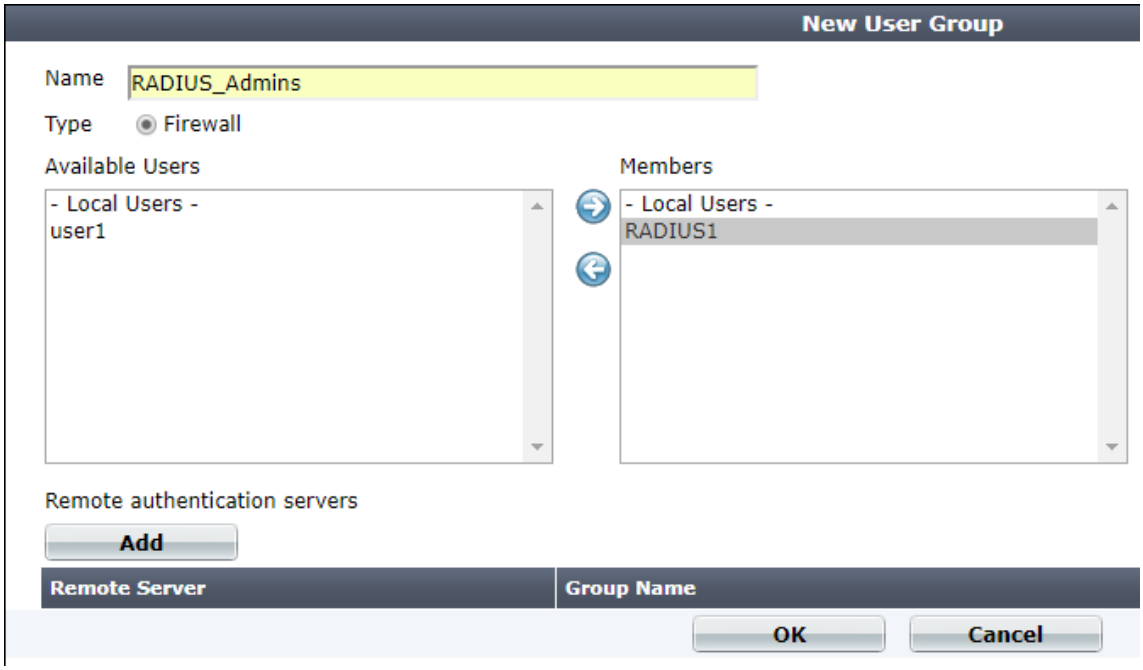

e. Click **OK**.

Using the CLI:

1. Create a RADIUS system admin group:

```
config system admin
   edit "RADIUS_Admins"
      set remote-auth enable
      set accprofile "super_admin"
      set wildcard enable
```

```
      set remote-group "RADIUS_Admins"
   next
end
```

2. Create a user:

```
config user radius
   edit "RADIUS1"
      set secret ENC
          6rF7O4/Zf3p2TutNyeSjPbQc73QrS21wNDmNXd/rg9k6nTR6yMhBRsJGpArhle6UOCb7b8InM3nrC
          euVETr/a02LpILmIltBq5sUMCNqbR6zp2fS3r35Eyd3IIrzmve4Vusi52c1MrCqVhzzy2EfxkBrx5
          FhcRQWxStvnVt4+dzLYbHZ
   next
end
```

3. Create a user group:

```
config user group
   edit "RADIUS_Admins"
      set member "RADIUS1"
   next
end
```

RADIUS accounting and FortiGate RADIUS single sign-on

NOTE: To obtain a valid Framed-IP-Address attribute value, you need to manually configure DHCP snooping in the 802.1x-authenticated ports of your VLAN network for both port and MAC modes.

You can use your FortiSwitch for RADIUS single sign-on (RSSO) in two modes:

- Standalone mode
- FortiLink mode (FortiSwitch managed by FortiGate)

FortiSwitch uses 802.1x-authenticated ports to send five types of RADIUS accounting messages to the RADIUS accounting server to support FortiGate RADIUS single sign-on:

- START—The FortiSwitch has been successfully authenticated, and the session has started.
- STOP—The FortiSwitch session has ended.
- INTERIM—Periodic messages sent based on the value set using the `set acct-interim-interval` command.
- ON—FortiSwitch will send this message when the switch is turned on.
- OFF—FortiSwitch will send this message when the switch is shut down.

Configuring the RADIUS accounting server and FortiGate RADIUS single sign-on

Use the following commands to set up RADIUS accounting and enable a FortiSwitch to receive CoA and disconnect messages from the RADIUS server:

```
config user radius
   edit <RADIUS_server_name>
      set acct-interim-interval <seconds>
      set secret <secret_key>
      set server <server_name_IPv4>
      config acct-server
         edit <entry_ID>
```

```
            set status {enable | disable}
            set server <server_IP_address>
            set secret <secret_key>
            set port <port_number>
            set source-ip <source_IP_address>
         next
      end
   next
end
```

Variable	Description
<RADIUS_server_name>	Enter the name of the RADIUS server that will be sending CoA and disconnect messages to the FortiSwitch. By default, the messages use port 3799.
acct-interim-interval <seconds>	Enter the number of seconds between each interim accounting message sent to the RADIUS server. The value range is 60-86400. The default is 600.
secret <secret_key>	Enter the shared secret key for authentication with the RADIUS server.
server <server_name_IPv4>	Enter the domain name or IPv4 address for the RADIUS server. There is no default.
<entry_ID>	Enter the entry identifier. The value range is 0-20.
status {enable | disable}	Enable or disable RADIUS accounting. The default is disable.
server <server_IP_address>	Enter the IPv4 address of the RADIUS server that will be receiving the accounting messages. There is no default value.
secret <secret_key>	Enter the shared secret key for the RADIUS accounting server.
port <port_number>	Enter the port number for the RADIUS accounting server to receive accounting messages from the FortiSwitch. The default is 1813.
source-ip <source_IP_address>	Enter the IPv4 address of the server that will be sending accounting messages. The default is 0.0.0.0.

Example: RADIUS accounting and single sign-on

Use the following commands to set up RADIUS accounting:

```
config user radius
   edit "local-RADIUS"
      set server 10.0.23.5
      set secret ENC
          LE8xetYYGiE0bkQpBDdH6acilwkYROCos7XK2q5cNPhu8sUDW9/fvkgE+fVURgZGEzTsndt41gb+K+zV
          9m+nXCnoUXqivzQdt1UNlMxgKXADnCpXuiY966aJsYigmW/AZ1IM5kweUxvuHK8eqJkkT0nl64c8DID/
          LMAcCTx6JMapRCBS
      set auth-type ms_chap_v2
      set acct-interim-interval 1200
```

```
        config acct-server
           edit 1
              set status enable
              set server 10.0.23.5
              set secret ENC
                  LE8xetYYGiE0bkQpBDdH6acilwkYROCos7XK2q5cNPhu8sUDW9/fvkgE+fVURgZGEzTsndt41gb
                  +K+zV9m+nXCnoUXqivzQdt1UNlMxgKXADnCpXuiY966aJsYigmW/AZ1IM5kweUxvuHK8eqJkkT0
                  nl64c8DID/LMAcCTx6JMapRCBS
              set port 1813
              set source-ip 10.105.142.19
           next
        end
     next
  end
```

RADIUS change of authorization (CoA)

NOTE: For increased security, each subnet interface that will be receiving CoA requests must be configured with the `set allowaccess radius-acct` command.

FortiSwitch supports two types of RADIUS messages:

- CoA messages to change session authorization attributes (such as data filters and the session-timeout setting) during an active session.
- Disconnect messages (DMs) to flush an existing session. For MAC-based authentication, all other sessions are unchanged, and the port stays up. For port-based authentication, only one session is deleted.

RADIUS CoA messages use the following Fortinet proprietary attribute:

```
Fortinet-Host-Port-AVPair 42 string
```

The format of the value is as follows:

Attribute	Value	Description
Fortinet-Host-Port-AVPair	action=bounce-port	The FortiSwitch disconnects all sessions on a port. The port goes down and then up again.
Fortinet-Host-Port-AVPair	action=disable-port	The FortiSwitch disconnects all session on a port. The port goes down until the user resets it.
Fortinet-Host-Port-AVPair	action=reauth-port	The FortiSwitch forces the reauthentication of the current session.
session-timeout	<session_timeout_value>	The FortiSwitch disconnects a session after the specified number of seconds of idleness. This value must be more than 30 seconds.

The RADIUS CoA server uses the following error codes in the disconnect messages:

Error Message	Error Code	Description
RADIUS_ERROR_CODE_UNSUPPORTED_ATTRIBUTE	401	The attribute is not supported.
RADIUS_ERROR_CODE_NAS_ID_MISMATCH	403	The FortiSwitch identification does not match the RADIUS identification.
RADIUS_ERROR_CODE_INVALID_ATTRIBUTE	407	The attribute is invalid.
RADIUS_ERROR_CODE_SESSION_NOT_FOUND	503	The session was not found.

Configuring CoA and disconnect messages

Use the following commands to enable a FortiSwitch to receive CoA and disconnect messages from a RADIUS server:

```
config system interface
   edit "mgmt"
      set ip <address> <netmask>
      set allowaccess <access_types>
      set type physical
      set snmp-index <index_number>
   next
config user radius
   edit <RADIUS_server_name>
      set radius-coa {enable | disable}
      set radius-port <port_number>
      set secret <secret_key>
      set server <server_name_IPv4>
   end
```

Variable	Description
ip <address> <netmask>	Enter the interface IP address and netmask.
allowaccess <access_types>	Enter the types of management access permitted on this interface. Valid types are as follows: `http https ping snmp ssh telnet radius-acct`. Separate each type with a space. You must include `radius-acct` to receive CoA and disconnect messages.
snmp-index <index_number>	Enter the SNMP index for this interface.
<RADIUS_server_name>	Enter the name of the RADIUS server that will be sending CoA and disconnect messages to the FortiSwitch. By default, the messages use port 3799.
radius-coa {enable | disable}	Enable or disable whether the FortiSwitch will accept CoA and disconnect messages. The default is disable.

Variable	Description
radius-port <port_number>	Enter the RADIUS port number. By default, the value is 1812.
secret <secret_key>	Enter the shared secret key for authentication with the RADIUS server.
server <server_name_IPv4>	Enter the domain name or IPv4 address for the RADIUS server. There is no default.

Example: RADIUS CoA

The following example enables the FortiSwitch to receive CoA and disconnect messages from the specified RADIUS server:

```
config system interface
   edit "mgmt"
      set ip 10.105.4.14 255.255.255.0
      set allowaccess ping https http ssh snmp telnet radius-acct
      set type physical
      set snmp-index 55
   next
config user radius
   edit "Radius-188-200"
      set radius-coa enable
      set radius-port 0
      set secret ENC
          +2NyBcp8JF3/OijWl/w5nOC++aDKQPWnlC8Ug2HKwn4RcmhqVYE+q07yI9eSDhtiIw63kR/oMBLGwFQo
          eZfOQWengIlGTb+YQo/lYJn1V3Nwp9sdkcblfyayfc9gTeqe+mFltKl5IWNI7WRYiJC8sxaF9Iyr2/l4
          hpCiVUMiPOU6fSrj
      set server "10.105.188.200"
   next
end
```

Viewing the CoA configuration

Use the following command to check the CoA settings:

```
S524DF4K15000024 # diagnose user radius coa

90075.874 DAS: :radius_das_diag_handler:
RADIUS DAS Server List:
radius2:
Type: RADIUS_8021X, IP: 10.105.252.79,
Last CoA/DM Client IP Addr    : 10.105.252.79
Disc Reqs     : 2
Disc ACKs     : 1
Disc NAKs     : 1
CoA  Reqs     : 0
CoA  ACKs     : 0
CoA  NAKs     : 0
radius3:
Type: RADIUS_8021X, IP: 10.105.252.76,
Last CoA/DM Client IP Addr    :
```

```
Disc Reqs      : 0
Disc ACKs      : 0
Disc NAKs      : 0
CoA  Reqs      : 0
CoA  ACKs      : 0
CoA  NAKs      : 0
```

Notes

- CoA and single sign-on are supported only by the CLI in this release.
- RADIUS CoA is supported only in standalone mode (not in FortiLink mode) in the current release.
- FortiSwitch supports using FortiAutheticator, FortiConnect, or Microsoft Network Policy Server (NPS) as the RADIUS server for CoA and RSSO.
- Each RADIUS server can support only one accounting manager in this release.
- RADIUS accounting is supported only when the eap-passthru mode is enabled.
- Fortinet recommends a unique secret key for each accounting server.
- For CoA to correctly function with FortiAutheticator or FortiConnect, you must include the User-Name and Framed-IP-Address attributes *or* the User-Name and Calling-Station-ID attributes in the CoA request.
- For the FSW-1048 platform, you must manually configure the port.
- To obtain a valid Framed-IP-Address attribute value, you need to manually configure DHCP snooping in the 802.1x-authenticated ports of your VLAN network for both port and MAC modes.
- MAB authentication does not support CoA/RSSO accounting messages.
- Statistics for RADIUS accounting messages are not supported in this release.
- By default, the accounting server is disabled. You must enable the accounting server with the `set status enable` command.
- The default port for FortiAuthenticator single sign-on is 1813 for FortiSwitch. The default port for the accounting CoA is 1646.

TACACS

This chapter contains information on using Terminal Access Controller Access-Control System (TACACS+) authentication with your FortiSwitch unit.

This chapter covers the following topics:

Administrative accounts

Administrative, or admin, accounts allow access to various aspects of the FortiSwitch configuration. The level of access is determined by the admin profile that is assigned to the admin account.

See Configuring administrator tasks on page 28 for the steps to create an admin profile.

Configuring a TACACS admin account

TACACS+ is a remote authentication protocol that provides access control for routers, network access servers, and other network computing devices using one or more centralized servers. If you have configured TACACS+ support and an administrator is required to authenticate using a TACACS+ server, the FortiSwitch contacts the TACACS+ server for authentication.

Using the Web-based manager:

1. Go to **System > Admin > Administrators** and select **Create New**.
2. Give the administrator account an appropriate name.
3. Set **Type** as **Remote.**
4. Set **User Group** to a group for remote users.
5. Enable **Wildcard**.
6. Set **Admin Profile** to use the new profile.
7. Select **OK**.

Using the CLI:

```
config system admin
   edit tacuser
      set remote-auth enable
      set wildcard enable
      set remote-group <group>
      set accprofile <profile>
   end
end
```

User accounts

User accounts identify a network user and determine what parts of the network the user is allowed to access.

Configuring a user account

```
config user tacacs+
   edit <tacserver>
      set authen-type {ascii | auto | chap | ms_chap | pap}
      set authorization enable
      set key <authorization_key>
      set server <server>
   end
end
```

Configuring a user group

```
config user group
   edit <tacgroup>
      set member <tacserver>
      config match
         edit 1
            set server-name <server>
            set group-name <group>
         end
      end
   end
end
```

Example configuration

The following is an example configuration of a TACACS+ user account, with the CLI syntax shown to create it:

1. Configuring a TACACS user account for login authentication:

```
config user tacacs+
   edit tacserver
   set authen-type ascii
   set authorization enable
   set key temporary
   set server tacacs_server
end
```

2. Configuring a TACACS+user group:

```
config user group
   edit tacgroup
      set member tacserver
      config match
         edit 1
            set server-name tacserver
            set group-name tacgroup
         end
```

```
      end
   end
end
```

3. Configuring a TACACS+ system admin user account:

```
config system admin
   edit tacuser
      set remote-auth enable
      set wildcard enable
      set remote-group tacgroup
      set accprofile noaccess
   end
end
```

Troubleshooting and support

FortiSwitch provides various features for troubleshooting and support.

This chapter covers the following topics:

Virtual wire

Some testing scenarios may require two ports to be wired 'back-to-back'. Instead of using a physical cable, you can configure a virtual wire between two ports. The virtual wire forwards traffic from one port to the other port with minimal filtering or modification of the packets.

Notes:

- ACL mirroring is not supported.
- You can select ports that are already ingress and egress mirror sources.

Using the Web-based manager:

1. Go to **Switch > Virtual Wire > Wire**.
2. Click **Create New** to create a new virtual wire.
3. Enter a name and select the ports for first member and second member.
4. Click **OK** to save the changes.

Using the CLI:

Use the following commands to configure a virtual wire:

```
config switch virtual-wire
   edit <virtual-wire-name>
      set first-member <port-name>
      set second-member <port-name>
      set vlan <vlan-id>
   next
end
```

Virtual wire ports set a special Tag Protocol Identifier (TPID) in the VLAN header. The default value is 0xdee5, a value that real network traffic never uses.

Use the following commands to configure a value for the TPID:

```
config switch global
   set virtual-wire-tpid <hex value from 0x0001 to 0xFFFE>
end
```

Use the following command to display the virtual wire configuration:

```
diagnose switch physical-ports virtual-wire list

  port1(1) to port2(2) TPID: 0xdee5 VLAN: 4011
  port3(3) to port4(4) TPID: 0xdee5 VLAN: 4011
  port5(5) to port25(25) TPID: 0xdee5 VLAN: 4011
  port7(7) to port8(8) TPID: 0xdee5 VLAN: 4011
```

Note the following information about virtual wire:

- Ports have ingress and egress VLAN filtering disabled. All traffic (including VLAN headers) is passed unchanged to the peer. All egress traffic is untagged.
- Ports have L2 learning disabled.
- Ports have their egress limited to their peer and do no allow egress from any other ports.
- The system uses TCAM to force forwarding from a port to its peer.
- The TCAM prevents any copy-to-cpu or packet drops.

TFTP network port

When you power on the FortiSwitch, the BIOS performs basic device initialization. When this activity is complete, and before the OS starts to boot, you can click any key to bring up the boot menu.

From the menu, click the "I" key to configure TFTP settings. With newer versions of the BIOS, you can specify the network port (where you have connected your network cable). If you are not prompted to specify the network port, you must connect your network cable to the default network port:

- If the switch model has a WAN port, the WAN port is the network port.
- If the switch has no WAN port, the highest port number is the network port.

Cable diagnostics

You can check the state of cables connected to a specific port. The following pair states are supported:

- Open
- Short
- Ok
- Open_Short
- Unknown
- Crosstalk

If no cable is connected to the specific port, the state is Open, and the cable length is 0 meters.

For supported models, see Supported models on page 11.

Using the Web-based manager:

1. Go to **Switch > Port > Physical**.
2. Select a port.
3. Click **Cable Diag**.
4. Click **Ok** to start the cable diagnostics.
   NOTE: Running cable diagnostics on a port that has the link up will interrupt the traffic for several seconds.
5. Click **OK** to close the Cable Diagnostics window.

Using the CLI:

Use the following command to run a time domain reflectometry (TDR) diagnostic test on cables connected to a specific port:

```
diagnose switch physical-ports cable-diag <physical port name>
```

NOTE: Running cable diagnostics on a port that has the link up will interrupt the traffic for several seconds.

For example:

```
# diagnose switch physical-ports cable-diag port1

port1: cable (4 pairs, length +/- 10 meters)
pair A Open, length 0 meters
pair B Open, length 0 meters
pair C Open, length 0 meters
pair D Open, length 0 meters
```

Use the following command to check the medium dependent interface crossover (MDI-X) interface status for a specific port:

```
diagnose switch physical-ports mdix-status <physical port name>
```

For example:

```
# diagnose switch physical-ports mdix-status port1

port1: MDIX(Crossover)
```

Selective packet sampling

NOTE: This feature is not supported on 3032.

During debugging, you might want to see whether a particular type of packet was received on an interface on the switch.

1. Set up an access control list (ACL) on the switch with the interface that you want to monitor. See Access control lists on page 104. This ACL is the ingress interface.
2. Set up a mirror for the “internal” interface.

For example, if you want to monitor interface port17 for any IP packet (ether-type 0x800) with a destination subnet of 10.10.10/24 and a source subnet of 20.20.20/24, use the following commands.

```
# show switch acl policy
config switch acl policy
```

```
    edit 1
       config action
          set mirror "internal"
       end
       config classifier
          set dst-ip-prefix 10.10.10.0 255.255.255.0
          set ether-type 0x0800
          set src-ip-prefix 20.20.20.0 255.255.255.0
       end
       set ingress-interface "port17"
    next
end
```

To examine the packets that have been sampled in the example, use the following command:

```
# diagnose sniffer packet sp17 none 6
```

Network monitoring

You can monitor specific unicast MAC addresses in directed mode, monitor all detected MAC addresses on a FortiSwitch in survey mode, or do both. The FortiSwitch gives the directed mode a higher priority than survey mode. The directed mode and survey mode are disabled by default.

NOTE: Network monitoring is not available on FSW-108D-POE or FSR-112D-POE.

Directed mode

In directed mode, you select which unicast MAC addresses that you want examined. The FortiSwitch detects various fields of the packet—such as MAC address, IP address, VLAN, and user name—and stores the data in either of two databases.

NOTE: You cannot specify broadcast or multicast MAC addresses.

The maximum number of MAC addresses that can be monitored depends on the FortiSwitch model.

Platform Series	Maximum Number of MAC Addresses Monitored	Maximum Number of Hosts
1xx, 2xx	10	250
4xx, 5xx	20	1,024
10xx, 30xx	30	4,096

To find out how many network monitors are available, use the following command:

```
diagnose switch network-monitor cfg-stats

Network Monitor Configuration Statistics:
---------------------------------
Adds         : 0
Deletes      : 0
Free Entries : 20
```

To find out which network monitors are being used currently, use the following command:

```
diagnose switch network-monitor dump-monitors
```

```
Entry ID       Monitor Type       Monitor MAC      Packet-count
================================================================
1               directed-mode   00:01:02:03:04:05       10
2               directed-mode   10:01:02:03:04:05       0
3               survey-mode     08:5b:0e:c1:07:65       419
4               survey-mode     08:5b:0e:4f:af:38       101
5               survey-mode     08:5b:0e:ce:59:40       2347
6               survey-mode     08:5b:0e:4f:af:44       0
7               survey-mode     08:5b:0e:c1:07:65       0
8               survey-mode     08:5b:0e:4f:af:38       80
9               survey-mode     08:5b:0e:ce:59:40       117
10              survey-mode     08:5b:0e:4f:af:44       0
```

To start network monitoring, use the following commands:

```
config switch network-monitor settings
   set status enable
end
```

To specify a single unicast MAC address (formatted like this: `xx:xx:xx:xx:xx:xx`) to be monitored, use the following commands:

```
config switch network-monitor directed
   edit <unused network monitor>
      set monitor-mac <MAC address>
   next
end
```

For example:

```
config switch network-monitor directed
   edit 1
      set monitor-mac 00:25:00:61:64:6d
   next
end
```

Survey mode

In survey mode, FortiSwitch detects MAC addresses to monitor for a specified number of seconds. You can specify network monitoring for 120 to 3,600 seconds. The default time is 120 seconds. The FortiSwitch detects various fields of the packet—such as MAC address, IP address, VLAN, and user name—and stores the data in either of two databases.

To start network monitoring in survey mode, use the following commands:

```
config switch network-monitor settings
   set status enable
   set survey-mode enable
   set survey-mode-interval <number of seconds>
end
```

For example:

```
config switch network-monitor settings
   set status enable
   set survey-mode enable
   set survey-mode-interval 480
end
```

Network monitoring statistics

After you have enabled network monitoring, you can view the statistics for the number and types of packets.

To see the type of packets going to and from monitored MAC addresses, use the following command:

```
diagnose switch network-monitor parser-stats

Network Monitor Parser Statistics:
----------------------------------
Arp         : 0
Ip          : 1
Udp         : 46
Tcp         : 353
Dhcp        : 0
Eapol       : 0
Unsupported : 352
```

To see the number of packets going to and from monitored MAC addresses, use the following command:

```
diagnose switch network-monitor dump-monitors

Entry ID       Monitor Type       Monitor MAC      Packet-count
=================================================================
1               directed-mode   00:01:02:03:04:05       10
2               directed-mode   10:01:02:03:04:05       0
3               survey-mode     08:5b:0e:c1:07:65       419
4               survey-mode     08:5b:0e:4f:af:38       101
5               survey-mode     08:5b:0e:ce:59:40       2347
6               survey-mode     08:5b:0e:4f:af:44       0
7               survey-mode     08:5b:0e:c1:07:65       0
8               survey-mode     08:5b:0e:4f:af:38       80
9               survey-mode     08:5b:0e:ce:59:40       117
10              survey-mode     08:5b:0e:4f:af:44       0
```

NOTE: The FortiSwitch creates an entry in the layer-3 database using the exact packet contents when they were parsed. If the MAC address is then assigned to a different VLAN, this change might not be detected immediately. If there is a discrepancy in the output for the `diagnose switch network-monitor dump-l2-db` and `diagnose switch network-monitor dump-l3-db` commands, use the output with the more recent time stamp.

To see all detected devices from the layer-2 database, use the following command:

```
diagnose switch network-monitor dump-l2-db

mac 00:01:02:03:04:05 vlan 1
created 19 secs ago, last seen 16 secs ago
```

```
user JoE sources: eapol
```

To see all detected devices from the IP address database, use the following command:

```
diagnose switch network-monitor dump-l3-db
```

```
mac 08:5b:0e:c1:07:65 ip 169.254.2.2 vlan 4094
created 63614 secs ago, last seen 2 secs ago
sources: arp ip
mac 00:10:20:30:40:50 ip 10.10.10.111 vlan 123
created 75 secs ago, last seen 45 secs ago
sources: arp ip
mac 00:11:22:33:44:55 ip 30.30.30.115 vlan 1
created 53 secs ago, last seen 53 secs ago
sources: dhcp arp ip
```

Deployment scenario

Working configuration for PC and phone for 802.1x authentication using MAC

Summary

A. Configure all devices.
   - PC
   - Phone
   - FortiSwitch
   - FortiAuthenticator
   - DHCP server

B. Authenticate phone using MAB and using LLDP-MED.

C. Authenticate PC using EAP 802.1x.

A. Configure all devices

1. Configure the PC, phone, FortiSwitch, FortiAuthenticator [RADIUS server], and DHCP server)

Phone configuration (file: macmode_phone_pc_ping_work)

1. On the phone, enable the WAN port and leave the VLAN ID at the default to allow LLDP-Med (Policy) designate for voice VLAN assignment.
2. On the phone, enable the LAN port and assign the VLAN ID for data matching the RADIUS VLAN assignment.

PC configuration

1. Install the supplicant software.
2. Launch the supplicant software, type the user name and password, and enable DHCP on the interface.

FortiSwitch configuration

1. Configure the LLDP profile for voice.

```
# show switch lldp
config switch lldp profile

edit "pexa" <<<<<<<<<<<<<<<<<
   set 802.1-tlvs port-vlan-id
      config med-network-policy
         edit "voice"
            set status enable
```

```
            set vlan 21
         next
         edit "voice-signaling"
            set status enable
            set vlan 31
         next
         edit "guest-voice"
         next
         edit "quest-voice-signaling"
         next
         edit "softphone-voice"
            set status enable
            set vlan 41
         next
         edit "video-conferencing"
         next
         edit "streaming-video"
         next
         edit "video-signaling"
         next
      end
   set med-tlvs inventory-management network-policy
```

2. Apply the LLDL profile on a dot1x port.

```
# show switch physical-port port4
config switch physical-port

   edit "pexa" <<<<<<<<<<<<<<<<
   set lldp-profile "pexa"
   set speed auto
   next
end
```

3. Configure a user group.

```
# show user group
config user group

   edit "Corp_Grp_10"
   set member "FAC_LAB"
   next
end
```

4. Configure the RADIUS server.

```
# show user radius
config user radius

   edit "FAC_LAB" <<<<<<<<<
   set secret

   ENCW82jBg06XhKD/4Dugqm8QF2f7D1B4bfFdDSZaLUQPwZXv4F8zMc5sWHRl9suwmbmzNnAnyqPaa
       rAYcSLuT8kVjFSRO0znx+TXVWTqdSeLCpbMv
       +HYFNOHMbYlfES8wTYYD40InCgrYr2johvr2vfa5KG4g8XMwKSIM0LurR//1WqT0fH
   set server
   next
```

```
end
```

5. Configure port security on the dot1x port.

   a. Configure mac-mode port-security.
   b. Add voice VLAN on allowed list (for example, 21).
   c. Apply the security group.

   Interface port4 configuration:

```
# show switch interface port4
config switch interface

   edit "port4"
   set allowed-vlans 20-21,31,41
   set security-groups "Corp_Grp_10"
   set snmp-index 4
configure port-security
   set auth-fail-vlan disable
   set guest-auth-delay 120
   set guest-vlan disable
   set mac-auth-bypass enable
   set port-security-mode 802.1X-mac-based
   set radius-timeout-overwrite disable
   set auth-fail-vlanid 40
   set guest-vlanid 30
end
```

RADIUS configuration

MAB Authentication:

- Add phone MAC address to MAB list.

802.1X Authentication

1. Create a local user.
2. Create a user group with "Attributes" and enable PEAP and MSChapv2.

DHCP configuration

1. On the DHCP server, configure a pool for phone and a pool for the PC.

```
!
ip dhcp pool PC
network 10.1.1.0 255.255.255.0
default-router 10.1.1.1
dns-server 10.1.1.1
!
ip dhcp pool PC
network 20.1.1.0 255.255.255.0
default-router 20.1.1.1
dns-server 20.1.1.5
```

2. Configure exclude lists for pools for both gateway and DNS.

```
ip dhcp excluded-address 20.1.1.1 20.1.1.1.5
```

```
<<<<gateway and dns server
ip dhcp excluded-address 10.1.1.1 10.1.1.1.5
<<<<gateway and dns server
!
ip dhcp pool PC
network 20.1.1.0 255.255.255.0
default-router 20.1.1.1
dns-server 20.1.1.5
```

3. Configure the switch port VLAN interface as a gateway for the phone.

```
# show run
Building configuration

Current configuration
!
interface vlan21 <<<<<<
ip address 20.1.1.1
end
```

4. Configure the switch port VLAN interface as a gateway for the PC.

```
# show run
Building configuration

Current configuration
!
interface vlan10 <<<<<<
ip address 10.1.1.1
end

#
```

5. Configure the l2 port and associate the voice VLAN.

```
# show run
Building configuration

Current configuration
!
interface GigabitEthernet g1/0/1 <<<<<<
switchport access vlan 21
switchport trunk encapsulation dot1q
switchport trunk all
switchport mode trunk
end
```

6. Configure the l2 port and associate the data VLAN.

```
# show run
Building configuration

Current configuration
!
interface GigabitEthernet g1/0/2 <<<<<<
switchport access vlan 10
```

```
switchport trunk encapsulation dot1q
switchport trunk all
switchport mode trunk
end
```

2. Connect a link between the FortiSwitch and the DHCP server and assign matching VLAN for the phone for both ports

3. Connect a link between the FortiSwitch and the DHCP server and assign a matching VLAN for the PC for both ports

B. Authenticate phone using MAB

1. Connect the phone to the switch to authenticate with RADIUS through the MAB (mac-bypass).
2. Once authenticated:
   a. On the FortiSwitch, verify that the port is authorized and that the voice VLAN is on the allowed list.

```
# diagnose switch 8 status
Signal 10 received - config reload scheduled

wrdapd_hostapd_dump_state_console Hostapd own address 90:6c:ac:18:6f:2f
dump_diag:1:
receive dump diagnostic 802_1x/MAB sessions. ifname :port4: dump_diag:1:

port4 : Mode: mac-based (mac-by-pass enable)
        Link: Link up
        Port State: authorized ( ) <<<<<<
        Native Vlan : 1
        Allowed Vlan list: 1,10,20-21,31,41 <<<<<<
        Untagged Vlan list:
        Guest VLAN:

        Client MAC Type Vlan Dynamic-Vlan
        68:f7:28:fb:c0:0f 802.1x 1 10
<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<phone

        Sessions info:
        68:f7:28:fb:c0:0f Type=802.1x,PEAP,state=AUTHENTICATED
        params:reAuth=3600
        00:a8:59:d8:f1:f6 Type=MAB,,state=AUTHENTICATED
        params: reAuth=3600

        edited on: 2016-11-29 17:25

        edited on: 2016-11-29 17:59
```

   b. On the FortiSwitch, verify that the lldp neighbor detail accurately reflects the phone and voice VLAN designation.

```
        Neighbor learned on port4 by LLDP protocol
Last change 140 seconds ago
Last packet received 13 seconds ago

Chassis ID: 20.1.1.10 (ip) <<<<<<<<<<
System Name: FON-670i
```

```
System Description
V12.740.335.12.B

Time To Live: 60 seconds
System Capabilities: BT
Enabled Capabilities: BT
MED type: Communication Device Endpoint (Class III)
MED Capabilities: CP
Management IP Address: 20.1.1.10

Port ID: 00:a8:59:d8:f1:f6 (mac) <<<<<<<<<<<<<<<
Port description: WAN Port 10M/100M/1000M
IEEE802.3, Power via MDI:
Power devicetype: PD
PSE MDI Power: Not Supported
PSE MDI Power Enabled: No
PSE Pair Selection: Can not be controlled
PSE power pairs: Signal
Power class: 1
Power type: 802.3at off
Power source: Unknown
Power priority: Unknown
Power requested: 0
Power allocated: 0
LLDP-MED, Network Policies:
voice: VLAN: 21 (tagged), Priority: 0 DSCP: 0 <<<<<<<<<<<<
voice-signaling: VLAN: 21 (tagged), Priority: 0 DSCP: 0
streaming-video: VLAN: 21 (tagged), Priority: 0 DSCP: 0

# Checking STA 00:a8:59:d8:f1:f6 inactivity:
Station has been active
```

c. On the phone, verify that the DHCP address is assigned.

d. On the DHCP server, check binding and ping from gateway to verify that the phone is reachable.

```
# show ip dhcp binding
IP address Client-ID/ Lease expiration Type
Hardware address
20.1.1.10 00a8.59d8.f1f6 Mar 20 1993 01:52 AM Automatic
#
#
#
# show ip dhcp binding
IP address Client-ID/ Lease expiration Type
Hardware address
10.1.1.7 0168.f728.fbc0.0f Mar 11 1993 01:54 AM Automatic <<<<<< pc
20.1.1.10 00a8.59d8.f1f6 Mar 20 1993 01:52 AM Automatic <<<<< phone
# ping 10.1.1.7

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.1.7, timeout is 2
!!!!!
seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/8 ms
# ping 10.1.1.7
```

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.1.7, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/8 ms
# ping 10.1.1.7

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.1.7, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/4/9 ms
# ping 20.1.1.10

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 20.1.1.10, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/8 ms
#
```

C. Authenticate the PC using EAP dot1x

1. Connect the PC to the phone for EAP authentication and VLAN assignment (for data)
2. Once authenticated:
   a. On the FortiSwitch, verify that the port is authorized and that the data VLAN assigned to dynamic has been placed on the allowed list.

```
# diagnose switch 8 status
Signal 10 received - config reload scheduled

wrdapd_hostapd_dump_state_console Hostapd own address 90:6c:ac:18:6f:2f
dump_diag:1:
receive dump diagnostic 802_1x/MAB sessions. ifname :port4: dump_diag:1:

port4 : Mode: mac-based (mac-by-pass enable)
        Link: Link up
        Port State: authorized ( ) <<<<<<
        Native Vlan : 1
        Allowed Vlan list: 1,10,20-21,31,41
   <<<<<<
        Untagged Vlan list:
        Guest VLAN:

        Client MAC Type Vlan Dynamic-Vlan
        68:f7:28:fb:c0:0f 802.1x 1 10
<<<<<<<<<<<<<<<<<<<<<<<< PC
        00:a8:59:d8:f1:f6 MAB 1 0

        Sessions info:
        68:f7:28:fb:c0:0f Type=802.1x,PEAP,state=AUTHENTICATED
        params:reAuth=3600
        00:a8:59:d8:f1:f6 Type=MAB,,state=AUTHENTICATED

        params:reAuth=3600

        edited on: 2016-11-29 17:25
```

```
edited on: 2016-11-29 17:59
```

b. On the PC, verify that the DHCP address is assigned.
c. From the DHCP server, check the binding and a ping from gateway to verify that the PC is reachable.